DAWN PUTNAM SESSION

Putnam 1995 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1995 — Paper A **5**
1995 A1 5
1995 A2 7
1995 A3 10
1995 A4 13
1995 A5 16
1995 A6 19

Putnam 1995 — Paper B **23**
1995 B1 23
1995 B2 25
1995 B3 28
1995 B4 31
1995 B5 33
1995 B6 36

Putnam 1995 — Paper A

1995 A1

1995 A1. Let S be a set of real numbers which is closed under multiplication (that is, if a and b are in S, then so is ab). Let T and U be disjoint subsets of S whose union is S. Given that the product of any *three* (not necessarily distinct) elements of T is in T and that the product of any three elements of U is in U, show that at least one of the two subsets T, U is closed under multiplication.

Fields: abstract algebra (closure properties), logic of proof by contradiction **Position:** A1 – opener, purely structural, no computation **Answer:** At least one of T, U must be closed under multiplication; the statement follows from a two-line contradiction argument.

Layer 1 THE WHAT — surface analysis

The natural attempt. The instinct on reading "closed under multiplication" twice – once for S, once (partially) for T and U via triples – is to try to prove the claim directly: pick $t_1, t_2 \in T$ and attempt to force $t_1 t_2 \in T$ using only the given triple-closure facts. You reach for the tools you have. Multiply by a third copy of t_1: $t_1^2 t_2 \in T$, since that is a product of three elements of T. True, but useless – it tells you something about a *different* product, not about $t_1 t_2$ itself. Try involving U: take $t_1, t_2 \in T$ and some $u \in U$; there is no hypothesis at all that constrains a mixed product like $t_1 t_2 u$, because both triple-closure conditions are same-set conditions. You can multiply three elements of T, or three elements of U, but the hypotheses are silent on anything that mixes the two sets. Worse, the claim is a disjunction – "at least one of T, U" – so trying to prove T is closed head-on is potentially chasing the wrong half: maybe in a given configuration it is U that is closed and T that isn't, and no amount of manipulating T-only triples will ever produce a counterexample to T's closure or a proof of it, because the truth of the statement doesn't commit to a side. The direct approach stalls at exactly this wall: every manipulation available from the two triple-closure hypotheses stays inside one set at a time, and the object you actually want to control – a two-element product – never gets pinned down by three-element facts about a single set in isolation.

The model chain. Suppose, for contradiction, that neither T nor U is closed under multiplication. Then there exist $t_1, t_2 \in T$ with $t_1 t_2 \in U$, and there exist $u_1, u_2 \in U$ with $u_1 u_2 \in T$. (These exist precisely because $S = T \cup U$ is a disjoint union: if $t_1 t_2 \notin T$ it must lie in U, and symmetrically for $u_1 u_2$.)

Now consider the single real number

$$P = t_1 t_2 u_1 u_2.$$

Group it two ways, using that real multiplication is commutative and associative:

$$P = (t_1 t_2) \cdot u_1 \cdot u_2, \qquad P = t_1 \cdot t_2 \cdot (u_1 u_2).$$

In the first grouping, $t_1 t_2 \in U$ by assumption, and $u_1, u_2 \in U$, so P is the product of three elements of U; by the hypothesis on U, $P \in U$.

In the second grouping, $t_1, t_2 \in T$, and $u_1u_2 \in T$ by assumption, so P is the product of three elements of T; by the hypothesis on T, $P \in T$.

So $P \in T$ and $P \in U$ simultaneously, contradicting that T and U are disjoint. Hence the assumption was false: at least one of T, U is closed under multiplication.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin at the same point – staring at the same product of three elements of T (or U) and asking what it buys you. They split at the decision of *how to attack a disjunction*. The natural attempt tries to prove one disjunct outright, using only the hypotheses on that one set, and finds there is nothing there to grab onto: three-fold closure of T alone cannot control a two-fold product without outside information. The model instead negates the whole disjunction – assumes *both* disjuncts fail – which is the move that manufactures the missing outside information: two concrete counterexample pairs, $t_1t_2 \in U$ and $u_1u_2 \in T$, that can now be fed into each other's hypotheses.

> *The claim "at least one of T, U is closed" is a disjunction, and I have no way to choose which one is true from the hypotheses alone – so instead of trying to prove either disjunct directly, I should negate the whole statement and assume* both *fail, because that hands me two concrete witnesses, $t_1t_2 \in U$ and $u_1u_2 \in T$, and concrete witnesses are things I can actually plug into the given closure hypotheses, unlike an abstract "T is closed" I'm trying to establish from nothing.*

> *The two triple-closure hypotheses are each blind to the other set – they never let a T-element and a U-element interact directly – so the only way to use both of them on the same information is to find a single product that has two different three-factor decompositions, one entirely inside T and one entirely inside U; multiplying my two witness products together, $t_1t_2 \cdot u_1u_2$, does exactly that, because real multiplication lets me regroup the same four factors as $(t_1t_2) \cdot u_1 \cdot u_2$ or as $t_1 \cdot t_2 \cdot (u_1u_2)$.*

Layer 3 THE WHY

Why does negating the disjunction, rather than attacking either half, count as more than a proof-technique reflex? Because of what the two closure hypotheses actually are: each one is a rule that only fires when you already have three elements from the *same* set. A direct proof of "T is closed" would have to start from an arbitrary pair $t_1, t_2 \in T$ and end at $t_1t_2 \in T$, and nothing in the problem supplies a bridge from two elements of T to a third element of T that you don't already have – the triple-closure hypothesis needs three inputs from T to produce one output in T, so it can never certify a bare product of two. The only way to generate a usable triple is to have a third factor already sitting in the right set, and the single richest source of "elements known to be in T" or "elements known to be in U" that the problem hands you for free is exactly a failure witness – a pair whose product lands in the *other* set, because that lets you assemble genuinely new triples.

This is also why negating *both* halves, not just one, is necessary: negating only "T is closed" gives you $t_1t_2 \in U$, one useful new element of U, but nothing symmetric to close the loop with – you'd have three elements of U available (t_1t_2, and you'd still need two more genuine elements of U to invoke U's triple-closure, which the problem does not hand you). It is only by additionally assuming U is not closed that you obtain $u_1, u_2 \in U$ with $u_1u_2 \in T$ – a second failure witness that supplies exactly the missing elements on both sides. The product $t_1t_2u_1u_2$ is then not a clever guess; it is the unique minimal object built entirely from the four elements the double negation hands you, and its two groupings are forced by nothing more than commutativity of real multiplication (associativity guarantees the groupings denote the same number in the first place). Disjointness of T, U is the axiom that turns "this number is in

both T and U" into an outright contradiction – it is not incidental to the argument, it is the entire reason a proof by contradiction closes here rather than merely producing a curious coincidence.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the conclusion is a disjunction – "at least one of T, U" – over hypotheses that are each same-set closure rules. On seeing it, write down: a disjunction with no way to decide sides from the hypotheses alone should be attacked by negating the whole thing, not by trying to prove either half outright.
2. That cue suggests proof by contradiction specifically because negating the disjunction produces two concrete failure witnesses (one pair from each set), and concrete witnesses are the only things the triple-closure hypotheses can actually consume.
3. Hesitation occurs right after obtaining $t_1t_2 \in U$ and $u_1u_2 \in T$: it is not obvious what to do with two unrelated-looking facts. It is reinforced away by noticing the hypotheses require exactly three same-set elements to fire, and counting that you now have, between the two witnesses, precisely four elements – enough to assemble two different triples if you multiply all four together and regroup.
4. The trap is treating the problem as "prove T is closed" (or "prove U is closed") as a direct target, spending effort trying to derive one specific disjunct instead of recognising that the statement is symmetric in T and U and the proof should be too.
5. Recognisable faster next time by naming the shape: two same-structure hypotheses on two disjoint pieces of a whole, with a disjunctive conclusion about a stronger version of that structure holding on at least one piece.
6. Reusable rule: to prove "at least one of two similarly-defined objects has property P" when P is itself definable via a counterexample (a pair whose product/result escapes the object), negate the whole disjunction to get one counterexample from each object, then look for a single combined construction (here, a product of all the witnesses) that can be decomposed in two different ways, each way certified by one of the two hypotheses – and let disjointness (or any two mutually exclusive containers) supply the final contradiction.

▷ **Cue:** prove "at least one of two disjoint, similarly-defined pieces of a whole has property P", where P's failure supplies a concrete witness $\longrightarrow$ **Move:** negate the whole disjunction to get one failure witness from each piece, then build a single combined object decomposable two ways, each decomposition certified by one piece's hypothesis, and let disjointness force a contradiction

1995 A2

1995 A2. For what pairs (a, b) of positive real numbers does the improper integral

$$\int_b^\infty \left(\sqrt{\sqrt{x+a} - \sqrt{x}} - \sqrt{\sqrt{x} - \sqrt{x-b}} \right) dx$$

converge?

Fields: real analysis (asymptotic expansion, comparison/p-test for improper integrals) **Position:** A2 – second problem, computational but requires a genuinely two-order Taylor expansion to close **Answer:** The integral converges if and only if $a = b$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The two radicals subtracted inside each square root are begging to be rationalised: $\sqrt{x+a} - \sqrt{x} = a/(\sqrt{x+a} + \sqrt{x})$, which for large x behaves like $a/(2\sqrt{x})$. Doing

the same for the second term, $\sqrt{x} - \sqrt{x-b} \sim b/(2\sqrt{x})$. Taking square roots of both leading approximations,

$$\sqrt{\sqrt{x+a} - \sqrt{x}} \sim \sqrt{\tfrac{a}{2}}\, x^{-1/4}, \qquad \sqrt{\sqrt{x} - \sqrt{x-b}} \sim \sqrt{\tfrac{b}{2}}\, x^{-1/4},$$

so the integrand behaves like $\left(\sqrt{a/2} - \sqrt{b/2}\right) x^{-1/4}$ for large x. Since $\int^{\infty} x^{-1/4}\, dx$ diverges ($p = 1/4 < 1$), any nonzero coefficient here forces divergence – so $a = b$ is necessary. This is real progress, and a student stops here, tempted to declare "hence the integral converges exactly when $a = b$," treating the vanishing of the leading term as self-evidently sufficient. But leading-order cancellation only proves the integrand is $o(x^{-1/4})$ – strictly smaller order – it says nothing about what that smaller order actually is. It could be $x^{-1/3}$ (still divergent, since $p = 1/3 < 1$), or $x^{-1} \log x$ (divergent), or anything at all; "smaller" is not "integrable." Declaring victory at this point is asserting a conclusion the work performed does not support – the wall here is invisible, which is what makes it dangerous: the argument feels finished because the numbers matched, but the actual convergence question (is the next-order term's exponent bigger than 1 in magnitude?) was never asked.

The model chain. Fix $a, b > 0$ and expand to *second* order, not first. Using $\sqrt{1+u} = 1 + \frac{u}{2} - \frac{u^2}{8} + O(u^3)$ for $u \to 0$, write, for large x,

$$\sqrt{x+a} = \sqrt{x}\left(1 + \frac{a}{2x} - \frac{a^2}{8x^2} + O(x^{-3})\right),$$

so

$$\sqrt{x+a} - \sqrt{x} = \frac{a}{2\sqrt{x}} - \frac{a^2}{8x^{3/2}} + O(x^{-5/2}) = \frac{a}{2\sqrt{x}}\left(1 - \frac{a}{4x} + O(x^{-2})\right).$$

Taking the square root of this (again to one further order, using $\sqrt{1+u} = 1 + u/2 + O(u^2)$):

$$\sqrt{\sqrt{x+a} - \sqrt{x}} = \sqrt{\frac{a}{2}}\, x^{-1/4}\left(1 - \frac{a}{8x} + O(x^{-2})\right) = \sqrt{\frac{a}{2}}\, x^{-1/4} - \frac{1}{8} a \sqrt{\frac{a}{2}}\, x^{-5/4} + O(x^{-9/4}).$$

By the identical computation with $-b$ in place of a (i.e. expanding $\sqrt{x} - \sqrt{x-b}$),

$$\sqrt{\sqrt{x} - \sqrt{x-b}} = \sqrt{\frac{b}{2}}\, x^{-1/4} + \frac{1}{8} b \sqrt{\frac{b}{2}}\, x^{-5/4} + O(x^{-9/4}).$$

Subtracting,

$$\sqrt{\sqrt{x+a} - \sqrt{x}} - \sqrt{\sqrt{x} - \sqrt{x-b}} = \left(\sqrt{\frac{a}{2}} - \sqrt{\frac{b}{2}}\right) x^{-1/4} - \frac{1}{8}\left(a\sqrt{\frac{a}{2}} + b\sqrt{\frac{b}{2}}\right) x^{-5/4} + O(x^{-9/4}).$$

Case $a \neq b$. The coefficient of $x^{-1/4}$ is nonzero, and $x^{-1/4}$ dominates every other term written, so the integrand is asymptotic to a nonzero constant times $x^{-1/4}$; since $\int_b^{\infty} x^{-1/4}\, dx$ diverges, so does the original integral.

Case $a = b$. The $x^{-1/4}$ coefficient vanishes exactly, leaving

$$\text{integrand} = -\frac{1}{4} a \sqrt{\frac{a}{2}}\, x^{-5/4} + O(x^{-9/4}),$$

and both $\int^{\infty} x^{-5/4}\, dx$ and $\int^{\infty} x^{-9/4}\, dx$ converge, since $5/4 > 1$ and $9/4 > 1$. The integrand is also continuous on $[b, \infty)$ – the only place either inner square root could vanish is at the lower endpoints $x = 0$ (irrelevant, since $x \geq b > 0$ there) and $x = b$ itself (where $\sqrt{x-b} = 0$, giving a finite value, not a singularity) – so there is no issue near $x = b$, only the tail at infinity matters. Hence the integral converges.

Combining the two cases: the integral converges if and only if $a = b$.

Caveat. The transcribed source solution's intermediate expansion line ("$\sqrt{x+a} - \sqrt{x} = x^{1/2}(1 + a/2x + O(x^{-2}))$") does not parse as a correct identity as written – it looks like an OCR or transcription artefact, since $\sqrt{x+a} - \sqrt{x}$ is $O(x^{-1/2})$, not $O(x^{1/2})$. The model chain above is my own independent derivation, carried to the order actually needed to settle convergence (the source's stated exponent $x^{-3/4}$ for the surviving term at $a = b$ does not match a term of that order in my expansion either, whose surviving leading term is order $x^{-5/4}$). Both derivations agree on the final answer, $a = b$, and my derivation is verified directly by two independent applications of the binomial series to a controlled order, so I present it as the model chain in place of reproducing the source's algebra verbatim.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree exactly through the first-order asymptotic $\left(\sqrt{a/2} - \sqrt{b/2}\right)x^{-1/4}$ and the conclusion that $a = b$ is *necessary*. They split immediately afterward: the natural attempt, having found that $a = b$ kills the leading term, declares the integral convergent "by symmetry" and stops. The model chain instead re-expands both radicals to *one further order* specifically to expose what power of x survives once the leading term cancels, and only then checks that surviving power against the p-test.

> *Showing the leading term vanishes when $a = b$ tells me the integrand is a strictly smaller order than $x^{-1/4}$ – but "smaller order" is not a number, and convergence of $\int^{\infty} x^{-p}\,dx$ depends on the exact value of p, not on whether p increased; I have not actually located p yet, so I cannot claim convergence until I expand one more term and read off what the surviving exponent is.*

> *Because the next term I need is the* second-*order term of a square-root expansion nested inside another square root, I need the binomial expansion $\sqrt{1+u} = 1 + u/2 + O(u^2)$ applied twice in succession – once to get $\sqrt{x+a} - \sqrt{x}$ to order $x^{-3/2}$, and again to take the square root of that whole expression to order $x^{-5/4}$ – rather than trusting that a single application (which only gave me the leading $x^{-1/4}$ term) has captured everything relevant.*

Layer 3 THE WHY

Why is "the leading term cancels" not the same fact as "the integral converges", even though both facts happen to have the same answer $a = b$ here? Because convergence of $\int^{\infty} f(x)\,dx$ for an f of definite eventual sign is entirely a statement about the *rate* of decay of f, and rate of decay is read off the exponent in $f(x) \sim Cx^{-p}$ – a fact that requires knowing p precisely, to a threshold of exactly 1. Cancelling the leading term only establishes an inequality ($f = o(x^{-1/4})$), and an inequality of that shape is compatible with both convergent behaviour (e.g. $p = 2$) and divergent behaviour (e.g. $p = 1/2$, or even $p = 1/4 - \varepsilon$ for tiny ε). The only way to pin down p exactly is to keep expanding the same series until a nonzero coefficient appears, because a Taylor/binomial expansion is, by definition, an ordered list of terms of strictly decreasing size – and "the size of the next term" is precisely what determines p once the first one vanishes.

This is also why the expansion has to be pushed through *two* layers of square root, not one: the quantity being expanded, $\sqrt{x+a} - \sqrt{x}$, is itself only known to two orders of accuracy in $1/x$ (order $x^{-1/2}$ then $x^{-3/2}$) from a single binomial expansion of $\sqrt{x+a}$; taking the square root of that two-term expansion is a second, independent application of the same binomial idea, and skipping it (stopping after only the outer square root's leading behaviour) is exactly what produces the false sense of completion in the natural attempt. The structural fact underneath all of this is the p-test itself: $\int^{\infty} x^{-p}dx$ converges iff $p > 1$, a hard threshold with no partial credit, so "the exponent got bigger" is meaningless until you know whether it crossed 1 – here it jumps from $p = 1/4$ (divergent) to $p = 5/4$ (convergent), and only an explicit second-order computation reveals that the jump clears the threshold rather than merely approaching it.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape "$\sqrt{\text{difference of two nearby square roots}}$" appearing twice with a subtraction between them; on seeing it, the move to write down immediately is "rationalise each inner difference, then Taylor-expand to leading order in $1/x$" to find the necessary condition on the parameters.

2. That cue suggests an asymptotic/comparison-test approach (rather than attempting to evaluate the integral in closed form), since the integrand's exact antiderivative is not realistically obtainable and only its tail behaviour is needed for a convergence question.

3. Hesitation occurs right at the moment the leading-order terms cancel: it feels like the problem is solved. It is reinforced away by asking explicitly, "what power of x survives now?" – a question that cannot be answered without redoing the expansion one order deeper, and by recalling that the p-test has a hard threshold at $p = 1$ that must be checked numerically, not just qualitatively.

4. The trap is mistaking "the obstruction to convergence (the leading term) is gone" for "the integral converges" – conflating a necessary condition, correctly derived, with a sufficient one that still needs its own argument.

5. Recognisable faster next time by treating every "does this parameterised tail integral converge" question as requiring the exponent of the surviving term to be pinned down exactly, and by expecting – whenever a leading term is designed to cancel – that the problem wants you to go one order further, since a competition would not set up a clean cancellation and then leave the next order irrelevant.

6. Reusable rule: when an asymptotic expansion of an integrand has a leading term that vanishes under some condition on the parameters, that condition is necessary but the argument is not finished – re-expand to the next nonzero order, identify its exact exponent, and check that exponent against the relevant convergence threshold ($p > 1$ for $\int^{\infty} x^{-p}dx$) before declaring convergence.

▷ **Cue:** difference of two similarly-shaped nested-radical expressions, integrated to infinity, with a "for what parameters does this converge" question $\longrightarrow$ **Move:** rationalise/Taylor-expand each term to leading order in $1/x$ to find the necessary cancellation condition on the parameters, then re-expand one order further under that condition and check the surviving exponent against the p-test threshold before claiming sufficiency

1995 A3

1995 A3. The number $d_1d_2\ldots d_9$ has nine (not necessarily distinct) decimal digits. The number $e_1e_2\ldots e_9$ is such that each of the nine 9-digit numbers formed by replacing just one of the digits d_i is $d_1d_2\ldots d_9$ by the corresponding digit e_i ($1 \leq i \leq 9$) is divisible by 7. The number $f_1f_2\ldots f_9$ is related to $e_1e_2\ldots e_9$ is the same way: that is, each of the nine numbers formed by replacing one of the e_i by the corresponding f_i is divisible by 7. Show that, for each i, $d_i - f_i$ is divisible by 7. [For example, if $d_1d_2\ldots d_9 = 199501996$, then e_6 may be 2 or 9, since 199502996 and 199509996 are multiples of 7.]

Fields: elementary number theory (modular arithmetic, place value) **Position:** A3 – third problem, notation-heavy but resolved by a single summation trick **Answer:** $d_i - f_i \equiv 0 \pmod 7$ for every i, proved by summing the nine per-position congruences to eliminate the position-dependent powers of 10.

Layer 1 THE WHAT — surface analysis

The natural attempt. The problem is stated entirely digit-by-digit – nine separate replacement conditions, one per position – so the natural first move is to work digit-by-digit too. Fix a

position i. Replacing d_i by e_i changes the number $D = d_1d_2\dots d_9$ by $(e_i - d_i)\cdot 10^{9-i}$, and the resulting number is divisible by 7, so

$$D + (e_i - d_i)\cdot 10^{9-i} \equiv 0 \pmod 7.$$

This pins down $e_i - d_i$ modulo 7, *provided* you know $10^{9-i} \bmod 7$ – which means computing the power-of-10 residues for $i = 1,\dots,9$, i.e. $10^8, 10^7, \dots, 10^0 \bmod 7$. Since $10 \equiv 3 \pmod 7$ and 3 has order 6 mod 7 $(3, 2, 6, 4, 5, 1, 3, \dots)$, this is a genuine, finishable computation: nine residues, repeating with period 6. Having found $10^{9-i} \bmod 7$ for each i, and knowing $D \bmod 7$ (which is not actually given!), you could in principle solve for $e_i - d_i \bmod 7$ position by position, and do the same again for $f_i - e_i$, then add the two to get $f_i - d_i$. But this runs straight into the missing piece: $D \bmod 7$ itself is never given, and there is no way to know it from nine separate replacement facts about individual digits without first establishing some relationship among D, E, and F as whole numbers – which is exactly the fact the position-by-position approach was trying to avoid computing. The digit-by-digit route is not wrong, but it is solving the same sub-problem (a relation between D and E as numbers) twice: once implicitly, buried inside nine separate modular-inverse computations, and it never gets there directly, because per-position algebra can only ever isolate $e_i - d_i$ in terms of the unknown $D \bmod 7$, not eliminate D altogether.

The model chain. Write $D = d_1d_2\dots d_9$, $E = e_1e_2\dots e_9$, $F = f_1f_2\dots f_9$ as integers (so $D = \sum_{i=1}^{9} d_i\cdot 10^{9-i}$, and similarly for E, F). The hypothesis says: for each $i = 1,\dots,9$, replacing d_i by e_i in D gives a multiple of 7, i.e.

$$D + (e_i - d_i)\cdot 10^{9-i} \equiv 0 \pmod 7 \qquad (i = 1,\dots,9). \tag{$*_i$}$$

Sum $(*_i)$ over all $i = 1,\dots,9$. The right-hand side is $9D$ plus $\sum_i (e_i - d_i)10^{9-i}$, and that sum telescopes: $\sum_i e_i 10^{9-i} - \sum_i d_i 10^{9-i} = E - D$ by definition of E and D. So

$$9D + (E - D) \equiv 0 \pmod 7 \quad\Longrightarrow\quad 8D + E \equiv 0 \pmod 7 \quad\Longrightarrow\quad D + E \equiv 0 \pmod 7,$$

using $8 \equiv 1 \pmod 7$ and $9 \equiv 2 \pmod 7$ carefully: precisely, $9D + (E - D) = 8D + E$, and $8 \equiv 1 \pmod 7$, giving $D + E \equiv 0 \pmod 7$.

The exact same argument, applied to the second batch of replacements (each of the nine numbers formed by replacing one e_i with f_i is divisible by 7), gives

$$E + F \equiv 0 \pmod 7.$$

Now fix any single position i and add $(*_i)$ to its analogue for the second batch,

$$E + (f_i - e_i)\cdot 10^{9-i} \equiv 0 \pmod 7, \tag{$**_i$}$$

to get

$$(f_i - d_i)\cdot 10^{9-i} + D + E \equiv 0 \pmod 7.$$

But $D + E \equiv 0 \pmod 7$ was just established, so this collapses to

$$(f_i - d_i)\cdot 10^{9-i} \equiv 0 \pmod 7.$$

Since $\gcd(10, 7) = 1$, 10^{9-i} is invertible modulo 7 (for every i, no case analysis on which power needed), so it cancels, leaving

$$f_i \equiv d_i \pmod 7, \quad \text{i.e.} \quad d_i - f_i \equiv 0 \pmod 7,$$

for every $i = 1,\dots,9$, exactly as required.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same nine congruences $(*_i)$, one per digit position. They split at the very next step: what to do with nine congruences that all share the unknown quantity D. The natural attempt tries to solve each $(*_i)$ individually for $e_i - d_i$, which requires knowing $10^{9-i} \bmod 7$ position by position and still leaves $D \bmod 7$ unknown. The model instead sums all nine congruences at once, which eliminates the position-dependent coefficients 10^{9-i} entirely (they telescope into $E - D$, a global fact) without computing a single power of 10 modulo 7.

> *Every one of these nine congruences names the same unknown, D, but with a different, position-dependent coefficient 10^{9-i} attached to the varying part – so instead of solving each one separately (which needs me to know both $10^{9-i} \bmod 7$ and $D \bmod 7$, and the second of those I don't have), I should add all nine together, because summing $(e_i - d_i) \cdot 10^{9-i}$ over every position i is, by the very definition of a 9-digit number as $\sum d_i 10^{9-i}$, exactly $E - D$: the position-dependent coefficients aren't eliminated by cancellation, they are absorbed into a fact about the whole numbers.*

> *Once I have $D + E \equiv 0 \pmod 7$ as a fact about the aggregate numbers, I can go back to a* single *position's congruence, add it to the corresponding congruence from the second replacement batch, and the two copies of the now-known quantity $D + E$ cancel out, leaving a clean statement purely about $(f_i - d_i) \cdot 10^{9-i} \bmod 7$ – and because 10 is coprime to 7, that power of ten is always invertible, so I never need its actual value.*

Layer 3 THE WHY

Why does summing the nine congruences work, rather than being merely a lucky shortcut? Because a 9-digit number is, by definition, $\sum_{i=1}^{9} (\text{digit}_i) \cdot 10^{9-i}$ – place value is itself a sum over positions – so any statement built by summing "one digit changed by δ_i" over all nine positions is automatically a statement about the change in the whole number, $\sum_i \delta_i \cdot 10^{9-i} = E - D$. The per-position congruences are not nine independent facts about nine independent unknowns; they are nine views of the same underlying object D, differing only in which digit is perturbed, and the problem is constructed so that perturbing every digit exactly once and adding the results reconstructs the entire number's transition from D to E. This is the real reason the position-by-position route is doing unnecessary work: it tries to recover $e_i - d_i$ digit by digit using the local coefficient 10^{9-i}, when the quantity the proof actually needs – $D + E \equiv 0 \pmod 7$ – lives one level up, at the level of the whole numbers, and is reached by exploiting the very definition of place value (a sum) rather than fighting it.

The second half of the WHY is why $D + E \equiv 0 \pmod 7$, once known, is enough to finish without ever touching a specific power of 10: subtracting two congruences that each contain $D + E$ (one from the $D \to E$ batch at position i, one from the $E \to F$ batch at the same position i) cancels the very quantity that required summing all nine relations to find, leaving only $(f_i - d_i) \cdot 10^{9-i} \equiv 0 \pmod 7$. Since 10 and 7 are coprime, multiplication by 10^{9-i} is a bijection on residues mod 7 for every single i – this is Bézout's identity in miniature, or just Fermat's little theorem giving $10^6 \equiv 1 \pmod 7$ so an explicit inverse exists – and a bijection can be cancelled from a congruence unconditionally, which is precisely why the proof never has to know which power of 10 is attached to which digit. The structural insight is that the problem hands you two "global" facts ($D + E \equiv 0$, $E + F \equiv 0$) for free once you sum correctly, and those two facts are exactly what is needed to make the local, per-digit statement collapse.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is nine congruences, each of the form (base number D) + (position-dependent term) $\equiv 0 \pmod 7$, all sharing the same base D. On seeing this, the move to write down is: sum all nine relations to see what happens to the shared term.

2. That cue suggests summation/elimination across the whole system of congruences, rather than solving each congruence separately for its position-dependent unknown.
3. Hesitation occurs at the point where the per-position approach requires knowing $10^{9-i} \bmod 7$ for all nine positions and, worse, $D \bmod 7$ itself, which is never given. It is reinforced away by noticing that the target fact – a relation between D and E as whole numbers – is not a per-digit fact at all, so a per-digit method was never going to reach it directly; summing is the operation that turns nine local facts into one global one.
4. The trap is computing the power-of-10 residues mod 7 for all nine positions (a real, absorbing distraction, since 10 has order 6 mod 7 and the exponents don't repeat evenly across 9 positions) in the belief that this granular information is what's needed, when the proof never actually requires knowing any individual $10^{9-i} \bmod 7$.
5. Recognisable faster next time by noticing that a "9-digit number" is defined as a sum over its digits with place-value weights, so any hypothesis given digit-by-digit, once summed over all digits, automatically reconstitutes a fact about the whole number – this pattern-match should fire immediately whenever a problem gives one relation per digit/position sharing a common base quantity.
6. Reusable rule: when a problem supplies several congruences (or equations) that each involve a shared "base" quantity plus a varying, position-indexed term, and the shared quantity itself is unknown, sum the relations across all positions before attempting to solve any single one – the varying terms telescope (by the very definition of the object, e.g. place value) into a fact about the aggregate quantities, eliminating the need to compute any individual position-dependent coefficient.

▷ **Cue:** several congruences/equations each sharing a common unknown base quantity plus a varying, position-indexed term (e.g. digit-by-digit perturbations of the same number) $\longrightarrow$ **Move:** sum the relations across all positions/indices so the varying terms telescope, by the object's own definition (e.g. place value), into a single fact about the aggregate quantities, before attempting to solve any individual relation

1995 A4

1995 A4. Suppose we have a necklace of n beads. Each bead is labeled with an integer and the sum of all these labels is $n-1$. Prove that we can cut the necklace to form a string whose consecutive labels $x_1, x_2, \ldots, x_n$ satisfy

$$\sum_{i=1}^{k} x_i \leq k-1 \qquad \text{for} \quad k = 1, 2, \ldots, n.$$

Fields: combinatorics (cyclic sequences, the cycle lemma); discrete extremal principle **Position:** A4 — middle of the paper; a short proof once the right quantity is maximised, and a long fight without it. **Answer:** Such a cut always exists: renormalise the partial sums so that one full lap adds nothing, and cut immediately after the bead where the renormalised partial sum is largest.

Layer 1 THE WHAT — surface analysis

The natural attempt. You look for a rule that picks the starting bead. The constraint at $k = 1$ says $x_1 \leq 0$, so the string must open on a bead with a non-positive label, and the obvious candidate is the smallest label on the necklace: start as low as possible and you have the most room later. Test it. Take $n = 5$ and the labels, in cyclic order,

$$-2,\ 6,\ -1,\ 0,\ 1 \qquad (\text{sum } 4 = n-1).$$

Starting at -2 gives partial sums $-2, 4, \ldots,$ and $4 > 1$ fails at $k = 2$. The smallest bead can be sitting right in front of the largest one.

So you try a repair rule: start somewhere; if the first failure is at position k, restart at bead $k+1$. This is the petrol-station argument and it can be made to work, but you now owe a proof that the restarts terminate — that every bead you skipped would also have failed, and that you never go round more than once. Each of those is a lemma with its own index bookkeeping, and under time pressure the bookkeeping is where the proof breaks. You also try induction on n by merging two adjacent beads a, b into one bead $a + b - 1$ (so the total becomes $n - 2$), but unmerging a good string of length $n - 1$ does not directly give a good string of length n: the constraint at the position between a and b is new and nothing controls it.

The wall is the same every time: you are choosing a bead by looking at *labels*, and the condition is about *partial sums*.

The model chain. Number the beads $1, \ldots, n$ round the necklace from an arbitrary starting point, with labels $y_1, \ldots, y_n$, and extend periodically: $y_{j+n} = y_j$. For $j \geq 0$ let

$$t_j = \sum_{i=1}^{j} y_i \; - \; j \cdot \frac{n-1}{n}, \qquad t_0 = 0.$$

One full lap adds $\sum_{i=1}^{n} y_i - n \cdot \frac{n-1}{n} = (n-1) - (n-1) = 0$, so $t_{j+n} = t_j$: the sequence (t_j) is periodic with period n.

Choose $m \in \{1, \ldots, n\}$ with $t_m = \max\{t_1, \ldots, t_n\}$; by periodicity $t_m \geq t_j$ for *every* $j \geq 0$. Cut between bead m and bead $m+1$, so the string is $x_k = y_{m+k}$ for $k = 1, \ldots, n$. For each k,

$$\sum_{i=1}^{k} x_i - k \cdot \frac{n-1}{n} = t_{m+k} - t_m \leq 0,$$

hence

$$\sum_{i=1}^{k} x_i \leq k \cdot \frac{n-1}{n} = k - \frac{k}{n} < k.$$

The left side is an integer strictly less than the integer k, so it is at most $k - 1$. □

(Check on the example: $t_0, \ldots, t_5 = 0,\ -2.8,\ 2.4,\ 0.6,\ -0.2,\ 0$, the maximum is at $m = 2$, and the string $-1, 0, 1, -2, 6$ has partial sums $-1, -1, 0, -2, 4$ against bounds $0, 1, 2, 3, 4$.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that the cut must be chosen deliberately. The natural attempt chooses by looking at individual labels (smallest bead, first failure). The model chooses by looking at the running total. The split is the question "which bead?" versus "which *prefix*?" — and, one step later, whether to subtract the average $\frac{n-1}{n}$ per bead before comparing prefixes.

> *The condition is about partial sums, not about labels. If I rotate the necklace, what happens to the partial sums? They all shift by the same amount — minus the partial sum at the cut. So I want to cut where the partial sum is as large as it gets; then everything after it is lower.*

> *But going round once adds $n - 1$, so the partial sums are not periodic and "the largest" is ambiguous. Subtract the average $\frac{n-1}{n}$ per bead; then a lap adds zero, the sequence is genuinely cyclic, and the maximum is a maximum over the whole necklace.*

> *The target is $k - 1$, but what I get is $k - \frac{k}{n}$. That is enough: all I need is "$< k$", and because the partial sum is an integer, $< k$ is the same as $\leq k - 1$.*

Layer 3 THE WHY

Why is "cut at the maximum partial sum" natural rather than a trick? Go back to what a rotation does. Fix a reference numbering and write P_j for the sum of the first j labels. If you cut after bead m, the new partial sums are $P_{m+k} - P_m$: every partial sum of the rotated string is a partial sum of the original, minus one fixed number P_m. So the whole family of rotations is a single sequence viewed from different starting heights. A condition of the form "every partial sum is at most something" is an upper bound on the graph of that sequence, measured from the starting point. The way to make every later value lie below your starting value is to start at the top. That is the extremal principle applied to the one object the problem is about, and it is forced as soon as you write the partial sums down.

The one wrinkle is that the raw partial sums drift: a lap adds $n-1$, so "the top" of a drifting sequence is not a well-defined place on a necklace. The drift is a linear function of j, so remove it by subtracting the average increment $\frac{n-1}{n}$ per bead. This is the same move as looking at a random walk relative to its mean, or at a function minus its linear part so that it becomes periodic. After it, t_j is a function on the necklace itself, and a function on a finite set has a maximum. Nothing is guessed: the correction is determined by the requirement "one lap adds zero".

Kedlaya's solution begins by setting $s_k = x_1 + \cdots + x_k - k(n-1)/n$. Here is where that comes from: it is the partial-sum sequence with its drift removed, which is the smallest change that makes "rotate the necklace" the same as "translate the sequence". His closing remark — that $i(n-1)/n$ "may be replaced by $i-1$ since the left side is an integer" — is carrying real weight. The fractional correction produces a bound $k - \frac{k}{n}$ that is not the bound asked for, and integrality is what converts it. This is also why the fractional version is cleaner than the integer version: if you subtract 1 per bead instead, the sequence falls by 1 per lap, you need strict inequalities, and you must pick the *last* index where the maximum is attained to get them. The fractional drift pays for strictness with integrality, so no tie-breaking is needed.

The problem's phrasing supplies every cue. "Necklace" and "cut" say: the answer is a rotation. "$\sum_{i=1}^{k} x_i \le \ldots$ for $k = 1, \ldots, n$" says: the object is the sequence of partial sums. "The sum of all these labels is $n-1$" says: one lap has a known total, which is exactly the datum you need to remove the drift. And "$k-1$" rather than "k" is the integer gap that the fractional correction falls into. This is the cycle lemma (Dvoretzky–Motzkin), which also counts ballot sequences and Catalan objects; the problem is a disguised form of it, and the necklace is the circle on which the lemma lives.

The natural attempt was reasonable: a greedy restart argument is a genuine proof strategy, and the petrol-station version of it is correct. It costs more because it proves the extremal statement indirectly, one failure at a time, instead of naming the extremal point directly.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the family of inequalities $\sum_{i \le k} x_i \le k-1$ for every k: a condition on all partial sums. On seeing it, write down the partial-sum sequence P_j and the formula "cutting after bead m gives partial sums $P_{m+k} - P_m$".

2. The extremal principle: cut at the maximum of the (drift-corrected) partial sums, so that every later partial sum is at most the starting one.

3. Hesitation comes when the raw partial sums are not periodic and "the maximum" seems ill-defined, or when the bound obtained is $k - \frac{k}{n}$ instead of $k-1$. Reinforce away the first by subtracting the average step so a lap adds zero; the second by recalling that an integer below k is at most $k-1$.

4. The trap is choosing the starting bead by its label (smallest label, first non-positive label). The counterexample $-2, 6, -1, 0, 1$ kills it in one line. A secondary trap is the integer drift correction without tie-breaking: with $P_j - j$ you must take the last maximiser, or equality at a later index breaks strictness.

5. Slogan: *a condition on all partial sums of a cyclic sequence is a condition on the height profile; rotate to start at the top.* And: *if the profile drifts, subtract the average step to make it periodic first.*

6. Decision rule: **when you may choose where a sequence starts and the constraint is on all its partial sums, remove the linear drift and start at the extremum of what remains; then use integrality to sharpen a fractional bound.**

▷ **Cue:** "cut the necklace" / "some rotation of a cyclic sequence satisfies …" ⟶ **Move:** write the rotated partial sums as $P_{m+k} - P_m$ and choose m extremal

▷ **Cue:** partial sums of a cyclic sequence whose full-lap total is $T \neq 0$ ⟶ **Move:** subtract jT/n so that the corrected sequence is periodic

▷ **Cue:** fractional bound $S < k$ (or $S \leq k - \frac{k}{n}$) on an integer quantity S, target $S \leq k - 1$ ⟶ **Move:** integrality: an integer strictly less than k is at most $k - 1$

▷ **Cue:** "all partial sums $\leq$ something" with free starting point (ballot, Catalan, petrol station) ⟶ **Move:** cycle lemma: exactly the rotations starting after a maximum work

1995 A5

1995 A5. Let $x_1, x_2, \ldots, x_n$ be differentiable (real-valued) functions of a single variable t which satisfy

$$\begin{aligned}
\frac{dx_1}{dt} &= a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\
\frac{dx_2}{dt} &= a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n \\
&\vdots \qquad\qquad\qquad\qquad\qquad\qquad \vdots \\
\frac{dx_n}{dt} &= a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nn}x_n
\end{aligned}$$

for some constants $a_{ij} > 0$. Suppose that for all i, $x_i(t) \to 0$ as $t \to \infty$. Are the functions $x_1, x_2, \ldots, x_n$ necessarily linearly dependent?

Fields: linear algebra; differential equations (linear systems with constant coefficients) **Position:** A5 — late in the paper, but the whole proof is three lines once the right scalar quantity is chosen; the difficulty is entirely in the choice. **Answer:** Yes. Some eigenvalue λ of $A = (a_{ij})$ has $\operatorname{Re}\lambda > 0$ (because $\operatorname{tr} A > 0$); for a left eigenvector w, the combination $w \cdot x$ satisfies $(w \cdot x)' = \lambda\,(w \cdot x)$, so it cannot tend to 0 unless it is identically 0, and that is a linear dependence.

Layer 1 THE WHAT — surface analysis

The natural attempt. Write the system as $\vec{x}\,' = A\vec{x}$. You know how to solve these: find the eigenvalues $\lambda_1, \ldots, \lambda_n$ and eigenvectors $\vec{v}_1, \ldots, \vec{v}_n$ of A, and the general solution is

$$\vec{x}(t) = \sum_{i=1}^{n} c_i e^{\lambda_i t}\vec{v}_i.$$

The functions $x_1, \ldots, x_n$ are the coordinates of this vector. For them to be linearly *independent* you need $\vec{x}(t)$ not to stay inside any proper subspace, which (you convince yourself) means every c_i must be non-zero. For every coordinate to tend to 0 you need every term present to decay, so every $\operatorname{Re}\lambda_i < 0$. So the question becomes: can a matrix with all entries positive have all its eigenvalues in the left half-plane?

That is the right question, but look at the cost of getting there. The general-solution formula assumes A is diagonalisable; if it is not, there are terms $t^k e^{\lambda t}$ and generalised eigenvectors, and the claim "independence forces every $c_i \neq 0$" has to be redone for Jordan blocks. Complex

eigenvalues produce $e^{\alpha t}\cos\beta t$ terms and you must pair them off to keep things real. Kedlaya's official solution goes this way, and even there the write-up has to wave at "everyone (presumably) knows" the structure of the solution space. It is a correct route through a lot of machinery, most of which the answer does not use.

The model chain. Answer: **yes**, they are necessarily linearly dependent.

Step 1: an eigenvalue in the right half-plane. The characteristic polynomial of A is $\det(\mu I - A) = \mu^n - (\operatorname{tr} A)\mu^{n-1} + \cdots$, and over $\mathbb{C}$ it factors as $\prod_{i=1}^{n}(\mu - \lambda_i)$. Comparing the μ^{n-1} coefficients, $\sum_i \lambda_i = \operatorname{tr} A = \sum_i a_{ii} > 0$. Taking real parts, $\sum_i \operatorname{Re}\lambda_i > 0$, so some eigenvalue λ has $\operatorname{Re}\lambda > 0$.

Step 2: a left eigenvector. Since $\det(A - \lambda I) = 0$ we also have $\det(A^T - \lambda I) = 0$, so there is a non-zero $w \in \mathbb{C}^n$ with $A^T w = \lambda w$, i.e. $\sum_i w_i a_{ij} = \lambda w_j$ for every j.

Step 3: one scalar equation. Put $y(t) = \sum_i w_i x_i(t)$, a differentiable complex-valued function. Then

$$y' = \sum_i w_i x_i' = \sum_i w_i \sum_j a_{ij} x_j = \sum_j \Big(\sum_i w_i a_{ij}\Big) x_j = \lambda \sum_j w_j x_j = \lambda y.$$

So $\big(e^{-\lambda t} y(t)\big)' = e^{-\lambda t}(y' - \lambda y) = 0$ (differentiate real and imaginary parts separately), and $y(t) = e^{\lambda t} y(0)$. Hence

$$|y(t)| = e^{(\operatorname{Re}\lambda) t}\,|y(0)| \;\geq\; |y(0)| \qquad (t \geq 0).$$

Step 4: decay forces vanishing. Each $x_i(t) \to 0$, so $y(t) \to 0$. With the inequality above, $|y(0)| \leq \lim |y(t)| = 0$, so $y(0) = 0$ and therefore $y \equiv 0$.

Step 5: a real relation. Write $w = u + iv$ with $u, v \in \mathbb{R}^n$. Since each x_i is real, $0 = \operatorname{Re} y = \sum_i u_i x_i$ and $0 = \operatorname{Im} y = \sum_i v_i x_i$ for all t. As $w \neq 0$, at least one of u, v is a non-zero real vector, and it gives constants, not all zero, with $\sum_i c_i x_i(t) = 0$ for every t. □

Example showing the situation is not vacuous. With $n = 2$ and $A = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ (eigenvalues $3, -1$), $\vec{x}(t) = e^{-t}(1, -1)$ is a solution tending to 0, and $x_1 + x_2 \equiv 0$; the coefficients $(1, 1)$ are the left eigenvector for $\lambda = 3$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from $\vec{x}\,' = A\vec{x}$. The natural attempt solves for $\vec{x}$ — all n coordinates at once, via the full eigen-decomposition. The model never solves the system; it picks a single linear combination $w \cdot \vec{x}$ whose evolution is a one-dimensional ODE. The split is at the first line: do you solve the system, or do you look for a combination that obeys a scalar equation of its own?

> *"Linearly dependent" means there are constants c_i with $\sum c_i x_i \equiv 0$. So what I am looking for is a vector c, and the object I should study is the scalar function $c \cdot \vec{x}(t)$. When does the ODE tell me something clean about $c \cdot \vec{x}$? When $c^T A$ is a multiple of c^T — a left eigenvector.*

> *Then $c \cdot \vec{x}$ is $e^{\lambda t}$ times a constant. If $\operatorname{Re}\lambda \geq 0$ that cannot go to zero unless the constant is zero. So I need one eigenvalue with non-negative real part — and positive entries give me positive trace, which is the sum of the eigenvalues.*

Layer 3 THE WHY

Why is the left eigenvector the natural object and not a trick? Start from the definition of the thing you are asked about. Functions $x_1, \ldots, x_n$ are linearly dependent if there is a constant vector $c \neq 0$ with $c \cdot \vec{x}(t) = 0$ for all t. So the conclusion is a statement about a *scalar* function $c \cdot \vec{x}(t)$, not about the vector $\vec{x}(t)$. The efficient plan is to study scalar functions of that form directly and ask which of them the differential equation controls. Differentiating,

$$(c \cdot \vec{x})' = c^T A \vec{x}.$$

This is again of the form (constant vector)$\cdot\vec{x}$, but with a different vector c^TA. The equation closes up on itself — becomes a scalar ODE in the single unknown $c\cdot\vec{x}$ — exactly when $c^TA=\lambda c^T$. That is the definition of a left eigenvector. Nothing is guessed: "I want a scalar combination whose derivative is a multiple of itself" *is* the eigen-equation for A^T.

Once you have $y=c\cdot\vec{x}$ with $y'=\lambda y$, the solution of a one-dimensional linear ODE is $y(0)e^{\lambda t}$, and the hypothesis $\vec{x}\to 0$ does the rest provided $|e^{\lambda t}|$ does not decay, i.e. $\operatorname{Re}\lambda\geq 0$. So the whole problem reduces to: *does A have an eigenvalue with non-negative real part?* The hypothesis $a_{ij}>0$ is there to answer that question, and the cheapest invariant that sees signs of eigenvalues is the trace. The diagonal entries alone already give $\operatorname{tr}A>0$; the off-diagonal positivity is not used. (Perron–Frobenius would give a positive real eigenvalue with a positive eigenvector, which is more than needed and far heavier.)

This also explains why the natural attempt is expensive. Solving $\vec{x}\,'=A\vec{x}$ means finding n independent quantities that each evolve by a scalar law — a full basis of (generalised) eigenvectors. The problem needs *one*. Kedlaya's solution constructs the full basis, then chooses a vector $\vec{w}$ orthogonal to all but one basis vector to isolate one coefficient; in the diagonalisable case that $\vec{w}$ is, up to scaling, a left eigenvector for the isolated eigenvalue, reached the long way round. The left eigenvector is the dual object — a coordinate function that the flow scales — and it exists as soon as one eigenvalue exists, with no diagonalisability hypothesis and no Jordan blocks.

There is also a structural reading of the answer. The set of initial conditions whose solutions tend to 0 is a subspace (the span of the generalised eigenspaces with $\operatorname{Re}\lambda<0$). The functions $x_1,\dots,x_n$ are dependent exactly when the solution lives in a proper A-invariant subspace. Positive trace forbids the stable subspace from being all of $\mathbb{R}^n$, so any decaying solution is trapped in a proper subspace; w is a linear functional that vanishes on that subspace. The proof above finds that functional without ever describing the subspace.

The question form "Are … necessarily …?" invites a counterexample hunt, and trying $n=2$ is sensible. Doing so is informative: with all entries positive you find you cannot make both eigenvalues negative (their sum is $a_{11}+a_{22}>0$), which is the trace argument in miniature and tells you the answer is yes.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "linearly dependent" applied to the solution of a linear system. On seeing it, write down: *I need a constant vector $c\neq 0$ with $c\cdot\vec{x}\equiv 0$; compute $(c\cdot\vec{x})'=c^TA\vec{x}$.*

2. Left eigenvector: choose c with $c^TA=\lambda c^T$, so $c\cdot\vec{x}$ satisfies a scalar ODE $y'=\lambda y$. Then locate an eigenvalue with $\operatorname{Re}\lambda\geq 0$ via the trace.

3. Hesitation occurs when λ and w are complex and you worry that "$y\equiv 0$" is not a real relation. Reinforce it away: the x_i are real, so the real and imaginary parts of w each give a real relation, and one of them is non-zero.

4. The trap is solving the whole system: diagonalisation, Jordan forms and complex conjugate pairing, all to extract one scalar fact. A second trap is reading "$a_{ij}>0$" as a cue for Perron–Frobenius when only the trace is needed.

5. Slogan: *to prove a linear relation among the coordinates of a solution, study the coordinate function the flow scales — a left eigenvector.* And: *signs of entries control eigenvalues through the trace first.*

6. Decision rule: **when the conclusion concerns one scalar combination of the unknowns, do not solve for all the unknowns; find the combination that the dynamics maps to a multiple of itself, and let the hypothesis act on that one scalar.**

▷ **Cue:** "are the solutions of $\vec{x}\,'=A\vec{x}$ linearly dependent" / a relation $c\cdot\vec{x}\equiv 0$ is wanted $\longrightarrow$ **Move:** differentiate $c\cdot\vec{x}$; choose c a left eigenvector so $(c\cdot\vec{x})'=\lambda(c\cdot\vec{x})$

▷ **Cue:** scalar $y'=\lambda y$ with $y\to 0$ and $\operatorname{Re}\lambda\geq 0$ $\longrightarrow$ **Move:** $|y(t)|=e^{(\operatorname{Re}\lambda)t}|y(0)|\geq|y(0)|$, so $y\equiv 0$

▷ **Cue:** sign information on the entries of a matrix, question about eigenvalues' real parts ⟶ **Move:** sum of eigenvalues = trace; positive trace forces some $\operatorname{Re}\lambda > 0$

▷ **Cue:** complex eigenvector giving a relation among real functions ⟶ **Move:** take real and imaginary parts; at least one is a non-zero real relation

1995 A6

1995 A6. Suppose that each of n people writes down the numbers 1,2,3 in random order in one column of a $3 \times n$ matrix, with all orders equally likely and with the orders for different columns independent of each other. Let the row sums a, b, c of the resulting matrix be rearranged (if necessary) so that $a \le b \le c$. Show that for some $n \ge 1995$, it is at least four times as likely that both $b = a+1$ and $c = a+2$ as that $a = b = c$.

Fields: probability (random walks on a lattice); combinatorics (multinomial coefficients); asymptotics
Position: A6 — the last problem of the morning paper, and it earns the position: exact computation is hopeless, and the proof needs two separate ideas (a one-step recursion between the two probabilities, and an exponential-versus-polynomial contradiction). **Answer:** Writing $C_n = \Pr(a = b = c)$ and A_n for the probability of one particular neighbouring configuration, $\Pr(b = a+1,\ c = a+2) = 6A_n$ and $C_{n+1} = A_n$; if $6A_n < 4C_n$ for all $n \ge 1995$ then A_n decays like $(2/3)^n$, but $A_{6m-1} \ge (6m+1)^{-5}$, a contradiction. (In fact the ratio tends to 6.)

Layer 1 THE WHAT — surface analysis

The natural attempt. You start by counting. The event $a = b = c$ means all three row sums equal $2n$; the other event means the row sums are $2n-1, 2n, 2n+1$ in some order. Each column is one of the 6 permutations of $(1,2,3)$, so if k_σ columns use permutation σ, the row sums are linear in the six numbers k_σ, and

$$\Pr(a = b = c) = \frac{1}{6^n} \sum \binom{n}{k_{\sigma_1}, \ldots, k_{\sigma_6}},$$

summed over all (k_σ) with $\sum k_\sigma = n$ that make the row sums equal. Working out the constraint: the first row sum is $k_{123} + k_{132} + 2k_{213} + \cdots$; setting all three equal gives two linear conditions, which with $\sum k_\sigma = n$ leave a three-parameter family of multinomials. The second event gives six such families. You can write these down. You cannot evaluate them at $n \ge 1995$, and comparing two three-fold sums of multinomials asymptotically means Stirling's formula inside a multi-dimensional sum — a local central limit theorem proved by hand.

Try small n instead. At $n = 2$, $\Pr(a = b = c) = \frac{1}{6}$ and the other event has probability $\frac{1}{3}$: a ratio of 2, not 4. So the statement is false at $n = 2$, which explains the strange "for some $n \ge 1995$": you are not being asked to prove an inequality for a given n, and the counting approach, which is built to handle a given n, is aimed at the wrong target.

The model chain. *Step 1: the walk.* Let R_1, R_2, R_3 be the row sums (unsorted) after n columns and put $D_n = (R_1 - 2n,\ R_2 - 2n,\ R_3 - 2n)$, a vector of integers with coordinate sum 0. A column is a uniformly random permutation of $(1,2,3)$; subtracting 2 from each entry, it adds to D a uniformly random element of

$$\Sigma = \{\text{the six permutations of } (1, 0, -1)\},$$

independently of earlier columns. So $D_0 = 0$ and D_n is a random walk with uniform steps from Σ.

Step 2: the two events. $a = b = c$ iff $R_1 = R_2 = R_3 = 2n$ iff $D_n = 0$. For the other event, sorted sums $a, a+1, a+2$ total $3a + 3 = 6n$, so $a = 2n - 1$ and $\{R_1, R_2, R_3\} = \{2n-1, 2n, 2n+1\}$, i.e. $D_n \in \Sigma$.

Step 3: symmetry. Permuting the three rows maps the step set Σ onto itself and preserves the uniform distribution on it, so it preserves the law of the whole walk: $\Pr(D_n = \pi v) = \Pr(D_n = v)$ for every coordinate permutation π. The set Σ is a single orbit under coordinate permutations, so $\Pr(D_n = s)$ is the same for all $s \in \Sigma$; call it A_n. Put $C_n = \Pr(D_n = 0)$. The claim is that

$$6A_n \;\geq\; 4C_n \qquad \text{for some } n \geq 1995.$$

Step 4: one step back from the origin. Condition on the last step. $D_{n+1} = 0$ iff $D_n = -s$ where s is the $(n+1)$st step. Since $-\Sigma = \Sigma$,

$$C_{n+1} = \sum_{s\in\Sigma} \tfrac{1}{6} \Pr(D_n = -s) = \sum_{s\in\Sigma} \tfrac{1}{6} A_n = A_n.$$

Step 5: suppose not. If $6A_n < 4C_n$ for every $n \geq 1995$, then by Step 4, $A_n < \frac{2}{3}C_n = \frac{2}{3}A_{n-1}$ for every $n \geq 1995$, so for $N > 1995$

$$A_N < \left(\tfrac{2}{3}\right)^{N-1995} A_{1995} \leq \left(\tfrac{2}{3}\right)^{N-1995}.$$

Step 6: a polynomial lower bound. Fix $m \geq 1$. If each of the six steps is used exactly m times in the first $6m$ columns, then $D_{6m} = 0$, because the six elements of Σ sum to 0 (each coordinate receives $1, 0, -1$ twice each). Hence

$$A_{6m-1} = C_{6m} \;\geq\; \frac{1}{6^{6m}} \cdot \frac{(6m)!}{(m!)^6}.$$

By the multinomial theorem, $6^{6m} = \sum \frac{(6m)!}{k_1!\cdots k_6!}$ over $k_1 + \cdots + k_6 = 6m$, $k_i \geq 0$. There are at most $(6m+1)^5$ terms ($k_1, \ldots, k_5$ each lie in $\{0, \ldots, 6m\}$ and determine k_6). The largest term is the balanced one: if $k_i \geq k_j + 2$, replacing (k_i, k_j) by $(k_i - 1, k_j + 1)$ multiplies the term by $\frac{k_i}{k_j+1} > 1$, so a maximising (k_i) has all entries within 1 of each other, and with sum $6m$ that forces every $k_i = m$. Therefore $\frac{(6m)!}{(m!)^6} \geq \frac{6^{6m}}{(6m+1)^5}$ and

$$A_{6m-1} \;\geq\; \frac{1}{(6m+1)^5}.$$

Step 7: contradiction. For $6m - 1 > 1995$, Steps 5 and 6 give $1 < (6m+1)^5 \left(\frac{2}{3}\right)^{6m-1996}$. The right side tends to 0 as $m \to \infty$ (an exponential with ratio $\left(\frac{2}{3}\right)^6 < 1$ beats the polynomial $(6m+1)^5$), so it is less than 1 for large m. Contradiction; hence $6A_n \geq 4C_n$ for some $n \geq 1995$. (The argument gives infinitely many such n.) □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains translate the events into "row sums equal" and "row sums consecutive". The natural attempt then tries to evaluate both probabilities. The model never evaluates either: at Step 4 it relates them by one step of the walk, $C_{n+1} = A_n$, which turns a comparison of two unknown quantities into a statement about consecutive terms of a single sequence. The split is at the question "how big are these?" versus "how are these two related?".

> *"For some $n \geq 1995$" — not "for all n", not "for $n = 1995$". I'm not being asked to compute anything. I'm being asked to show that a bad inequality can't hold forever. So: suppose it holds for every $n \geq 1995$ and find what grows or shrinks too fast.*

> *All-equal row sums is the origin; consecutive row sums is one step away from the origin. To be at the origin at time $n+1$ I must be one step away at time n. So $C_{n+1} = A_n$, and "A_n is small compared with C_n" is the same as "A_n is small compared with A_{n-1}".*

> *If $A_n / A_{n-1} < \frac{2}{3}$ forever, A_n dies exponentially. But a walk with no drift can't be exponentially unlikely to be home: the path using every step equally often already has probability like $1/\mathit{poly}(n)$.*

Layer 3 THE WHY

Why is the random-walk translation the natural first move, and not a piece of exotic apparatus? Because the matrix is built column by column, with independent columns, and the events concern only the row sums. Row sums are cumulative totals of independent, identically distributed column vectors: that is the definition of a random walk. Subtracting the mean column $(2,2,2)$ is the same drift-removal as in any walk — it puts the "typical" position at the origin and makes both events statements about positions near the origin. Once there, the two events are "at the origin" and "at one of the six neighbours", and the six neighbours are exactly the step set. That coincidence is not luck: "$b = a+1$, $c = a+2$" is the smallest non-trivial deviation from $a = b = c$, and the smallest deviation of a walk is one step.

Why relate C_{n+1} to A_n? Because the only thing you know about a random walk exactly is how it moves in one step. Any exact identity between probabilities at different positions must come from that. "To be home now I was one step away a moment ago" is the first identity the step structure gives, and here it happens to connect the two quantities the problem cares about. Kedlaya's solution states $C_{n+1} = A_n$ in a parenthesis; this is where it comes from.

Why does the contradiction work? Assuming the inequality fails from 1995 on gives a ratio bound $A_n / A_{n-1} < \frac{2}{3}$ at every step, and a ratio bound iterated is an exponential bound. On the other side, the walk has mean-zero steps, so it does not run away; the probability of being near home after n steps shrinks only like a power of n. The cheapest certificate of that is one explicit family of paths that return home: use each of the six steps equally often. The multinomial count of those paths is the largest term in the expansion of 6^{6m}, and a largest term among polynomially many terms is at least the total divided by the number of terms. That is pigeonhole, not Stirling. Kedlaya invokes Stirling to get the sharp rate $m^{-5/2}$; the argument needs only *some* polynomial, and $(6m+1)^{-5}$ is enough.

Why "$\geq$ 1995" and "four times"? The year is decoration: the proof works for any threshold, because the contradiction lives at $m \to \infty$. The factor 4 is slack. Numerically $6A_n/C_n$ is 5.46 at $n = 10$, 5.72 at $n = 20$ and 5.86 at $n = 40$, and it tends to 6 (all positions near the origin become equally likely — a local limit theorem, which I have not proved here and which the proof above does not use). What the proof does show is that A_n / A_{n-1} cannot stay below any fixed constant less than 1, so the ratio $6A_n/C_n = 6A_n/A_{n-1}$ exceeds $6(1-\varepsilon)$ infinitely often; $4 = 6 \cdot \frac{2}{3}$ is one convenient instance. Read the "for some" accordingly: proving "for all $n \geq 1995$" would need something like the local limit theorem, while "for some" needs only the two facts above.

Finally, the trap at small n is real: at $n = 2$ the ratio is 2. Any attempt that tries to establish the inequality monotonically from small cases is doomed, and the phrasing is warning you.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "each of n people … independent of each other" together with "row sums": a sum of independent identical steps. Write down $D_n = (R_1 - 2n, R_2 - 2n, R_3 - 2n)$ and the step set Σ. The second cue is "for some $n \geq 1995$": write down "suppose not, for all $n \geq 1995$".

2. Random walk plus last-step decomposition to get $C_{n+1} = A_n$; then contradiction by comparing a forced exponential decay with a polynomial lower bound from a single balanced multinomial.

3. Hesitation occurs at the lower bound: without Stirling, how small is $(6m)!/((m!)^6 6^{6m})$? Reinforce it away with the pigeonhole bound: it is the largest of at most $(6m+1)^5$ terms summing to 6^{6m}.

4. The trap is computing: the exact probabilities are multi-parameter sums of multinomials, and small cases ($n = 2$ gives ratio 2) actively mislead. A second trap is forgetting the symmetry that makes all six neighbours equally likely, which is needed to write the second event as $6A_n$.

5. Slogan: *"for some n beyond a threshold" means: assume it fails for all large n and find incompatible growth rates.* And: *a walk's probabilities at adjacent positions are linked by the last step.*

6. Decision rule: **when asked to compare two probabilities of a process built from independent steps, first relate them by one step of the process rather than computing either; if the conclusion is only "for some large n", turn a failed ratio into geometric decay and refute it with one explicit family of outcomes of polynomial probability.**

▷ **Cue:** a matrix or configuration built from independent identical columns, question about row totals $\longrightarrow$ **Move:** the totals form a random walk; subtract the mean step so the typical position is the origin

▷ **Cue:** two events "at the origin" and "one step from the origin" $\longrightarrow$ **Move:** condition on the last step: $\Pr(\text{home at } n+1) = \Pr(\text{a given neighbour at } n)$ when the steps are symmetric

▷ **Cue:** "show that for some $n \geq N$, $u_n \geq c\,u_{n-1}$" $\longrightarrow$ **Move:** suppose not; then $u_n < c^{n-N}u_N$, and refute with a lower bound decaying slower than c^n

▷ **Cue:** lower bound on a central multinomial $\binom{km}{m,\ldots,m}/k^{km}$ without Stirling $\longrightarrow$ **Move:** it is the largest of at most $(km+1)^{k-1}$ terms summing to k^{km}

▷ **Cue:** probability of the same position for all permutations of coordinates, symmetric step law $\longrightarrow$ **Move:** orbit symmetry: all positions in one orbit are equally likely

Putnam 1995 — Paper B

1995 B1

1995 B1. For a partition π of $\{1,2,3,4,5,6,7,8,9\}$, let $\pi(x)$ be the number of elements in the part containing x. Prove that for any two partitions π and π', there are two distinct numbers x and y in $\{1,2,3,4,5,6,7,8,9\}$ such that $\pi(x) = \pi(y)$ and $\pi'(x) = \pi'(y)$. [A *partition* of a set S is a collection of disjoint subsets (parts) whose union is S.]

Fields: combinatorics, pigeonhole principle, double counting **Position:** B1 — the opener of the second half **Answer:** Such x, y always exist. Proof: the number of distinct values $\pi(x)$ can take is at most 3, because distinct positive block sizes summing to 9 satisfy $1+2+3 = 6 \le 9 < 10 = 1+2+3+4$. If the claim failed, the map $x \mapsto (\pi(x), \pi'(x))$ would be injective into a grid of at most $3 \times 3 = 9$ pairs, forcing it to be a bijection onto all 9, which forces every value of $\pi(x)$ to occur exactly 3 times — but the count of x with $\pi(x) = s$ is always a multiple of s, so $s \mid 3$, leaving only $s \in \{1,3\}$ for what must be 3 distinct values. Contradiction.

Layer 1 THE WHAT — surface analysis

The natural attempt. The problem is begging for pigeonhole: nine elements, and to each one you attach the pair $(\pi(x), \pi'(x))$. If two elements ever get the same pair, you're done, so assume for contradiction that all nine pairs are distinct and try to find the contradiction there. The honest first move is to count how many pairs are even available. $\pi(x)$ is a block size, so naively it ranges over $1, \dots, 9$; same for $\pi'(x)$. That gives $9 \times 9 = 81$ possible pairs for only 9 elements — pigeonhole in the crude sense says nothing, since $9 \ll 81$.

At this point a strong student does not give up on pigeonhole, but doesn't yet see how to shrink the codomain either, so the natural next move is experimental: build small examples. Try $\pi =$ the all-singletons partition (every $\pi(x) = 1$) against $\pi' =$ the partition into one block of 9 ($\pi'(x) = 9$ for all x). Every pair is $(1, 9)$ — collision immediate, no content. Try something less degenerate: π splits $\{1, \dots, 9\}$ into blocks of sizes $2, 3, 4$ (so $\pi(x) \in \{2,3,4\}$), and π' into blocks of sizes $1, 3, 5$ (so $\pi'(x) \in \{1,3,5\}$). Now there really are up to $3 \times 3 = 9$ potential pairs, matching the 9 elements exactly — tantalisingly close to a genuine test case. Assign the blocks by hand, chasing an assignment where all nine pairs come out distinct, and it keeps failing: whichever block of size 3 in π you try to pair against blocks of π', the size-3 block always contributes 3 elements needing 3 distinct π'-values from that one block, and the arithmetic of which π'-block each of those 3 elements can fall into refuses to spread them out for every possible assignment. That failure is suggestive but it's a search over finitely many hand-built partitions, not a proof — and there is no guarantee the pattern holds once π or π' has a different size profile (say, blocks $1, 2, 6$). Without a general argument this is a dead end dressed as progress: hours of casework with no stopping point, and the crude 9 vs 81 pigeonhole never actually got used.

The model chain. Fix a partition π of a 9-element set. The values $\pi(x)$ takes are exactly the distinct block sizes present in π. Suppose k distinct sizes occur; since sizes are distinct positive integers, they sum to at least $1 + 2 + \cdots + k = \frac{k(k+1)}{2}$, and that sum is at most the total number

of elements, 9. Since $\frac{3\cdot4}{2} = 6 \leq 9$ but $\frac{4\cdot5}{2} = 10 > 9$, we get $k \leq 3$: $\pi(x)$ takes at most 3 distinct values. The identical argument bounds $\pi'(x)$ by at most 3 distinct values.

Now suppose, for contradiction, that no two distinct x, y satisfy $\pi(x) = \pi(y)$ and $\pi'(x) = \pi'(y)$; that is, the map $x \mapsto (\pi(x), \pi'(x))$ is injective on the 9-element ground set. Its image lies in a grid of at most $3 \times 3 = 9$ pairs. An injective map from a 9-element set into a set of size at most 9 forces that target set to have size at least 9 — so it has size exactly 9, meaning $\pi(x)$ really does take 3 distinct values, $\pi'(x)$ really does take 3 distinct values, and the map hits every one of the 9 cells of the 3×3 grid exactly once.

In particular, fix one value s that $\pi(x)$ attains. Because the map is a bijection onto the full grid, s is paired with each of the 3 values of $\pi'(x)$ exactly once — so exactly 3 elements x satisfy $\pi(x) = s$.

But the count of elements with $\pi(x) = s$ has a second, independent description: $\pi(x) = s$ exactly for the elements lying in a block of π of size s, and if there are m such blocks, that count is $m \cdot s$ — always a multiple of s. Setting $ms = 3$ forces $s \mid 3$, so $s \in \{1, 3\}$. This must hold for *every* one of the 3 distinct values $\pi(x)$ takes — but $\{1, 3\}$ has only two elements, and three distinct values cannot all lie in a two-element set. Contradiction. Hence the assumed injective map cannot exist, and the required x, y do.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same opening observation almost immediately: raw pigeolhole on 9 elements against $9 \times 9 = 81$ possible pairs proves nothing, so the codomain has to be shrunk before pigeonhole can bite. They split at *how* to shrink it. The natural attempt shrinks it case by case, by hand-picking specific block-size profiles and searching for a collision; the model chain shrinks it once and for all, by proving the general bound "at most 3 distinct values of $\pi(x)$, for any partition of a 9-set," from the arithmetic fact $1+2+3 \leq 9 < 1+2+3+4$.

> *I keep building specific examples with 3 block sizes on each side because that's the case where the pigeonhole grid, 9 cells for 9 elements, is exactly tight enough to be interesting — so instead of testing more examples, I should prove that 3 is the* maximum *number of distinct block sizes any partition of 9 elements can have, using nothing about which sizes, only that they're distinct positive integers summing to at most 9. If that bound holds for every partition, I don't need to search for a bad pairing at all: a 3×3 grid is the worst case there is, and I can attack that one case with full generality instead of attacking a hundred instances of it by hand.*

Layer 3 THE WHY

Why does bounding the number of distinct block sizes settle the whole problem, rather than just producing one more (larger) case to search by hand? Because it converts an open-ended search over partitions into a closed, finite, and in fact single case: however π and π' are chosen, the map $x \mapsto (\pi(x), \pi'(x))$ always lands in a grid no bigger than 3×3. That single arithmetic fact — k distinct positive integers sum to at least $k(k+1)/2$, the sum of the k smallest possible values — is a direct consequence of what distinctness among positive integers means, so the moment you notice the sizes are distinct (which is definitional: $\pi(x)$ is literally naming which block x is in, and a partition's blocks are disjoint), the bound $k \leq 3$ is forced, not guessed.

The second piece of the model chain — that the count of x with $\pi(x) = s$ is a multiple of s — is equally definitional rather than clever: $\pi(x) = s$ singles out exactly the elements sitting inside a block of size s, and blocks are disjoint by the definition given in the problem itself, so this count is literally s times (the number of blocks of that size). Nothing about the value of s or the shape of the partition is used beyond the definition of π.

What makes the argument close, and not merely constrain, is that these two definitional facts are forced to collide: the pigeonhole step (assume injectivity, get a 3×3 bijection) manufactures a very specific number — every value must occur exactly 3 times — and the divisibility fact

($s \mid \text{count}$) turns that specific number into a constraint on s itself ($s \mid 3$). A contradiction only appears because $\{1, 3\}$, the full set of divisors of 3 among positive integers, is too small to host the 3 distinct values that the bijection independently demanded. Each half of the argument is elementary; the archive's value here is noticing that the *same* number, 3, appears on both sides for two unrelated reasons (grid size from the size-bound, and the multiplicity from the divisibility fact) and that forcing them to agree is what kills the counterexample.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "partition of a 9-element set, and $\pi(x)$ names a block size." On seeing it, write down immediately: distinct positive block sizes present in π sum to at most 9, so at most 3 distinct sizes are possible ($1 + 2 + 3 = 6 \leq 9 < 10 = 1 + 2 + 3 + 4$).

2. That cue suggests bounding the range of the invariant before attempting any pigeonhole argument on the pair $(\pi(x), \pi'(x))$ — shrink the codomain from 9×9 down to at most 3×3, then apply pigeonhole (or, as here, assume injectivity and look for a contradiction) on the small grid.

3. Hesitation occurs right after noticing 9 vs 81 pigeonhole is useless: the instinct is to search hand-built examples for a collision rather than hunt for a hidden structural bound. This is reinforced away by asking, explicitly, "what have I not yet used about partitions of a fixed finite set?" — here, that block sizes are distinct positive integers with a bounded sum.

4. The trap is treating $\pi(x)$'s range as "up to 9" (one value per possible block size) rather than "at most 3" (the number of *distinct* sizes a single partition of 9 elements can actually contain) — the naive range is correct but nine times too generous to be useful.

5. Recognise this family faster by treating any invariant of the form "size of the part/block/group containing x" on a fixed finite ground set as automatically bounded in its number of distinct values by the largest k with $1 + 2 + \cdots + k \leq |S|$ — write that bound down before doing anything else with the invariant.

6. Reusable rule: when two invariants derived from partitions (or similar structures with bounded "distinct-value budgets") are combined via pigeonhole, first bound each invariant's range using the arithmetic of distinct positive parts summing to a fixed total, then look for a second, independent constraint on the same quantity (here, a divisibility fact from the block structure) to force the two bounds to collide.

▷ **Cue:** an invariant $\pi(x)$ equals the size of the block/part/group containing x, on a ground set of fixed size n $\longrightarrow$ **Move:** bound the number of distinct values $\pi(x)$ can take by the largest k with $1 + 2 + \cdots + k \leq n$, before attempting pigeonhole on $\pi(x)$ directly

▷ **Cue:** pigeonhole (or an injectivity assumption for contradiction) forces some value of a block-size invariant to occur exactly m times $\longrightarrow$ **Move:** check that m is compatible with the definitional fact that the count of elements with block size s is always a multiple of s — this often supplies the closing contradiction for free

1995 B2

1995 B2. An ellipse, whose semi-axes have lengths a and b, rolls without slipping on the curve $y = c \sin\left(\frac{x}{a}\right)$. How are a, b, c related, given that the ellipse completes one revolution when it traverses one period of the curve?

Fields: calculus (arc length), curve geometry **Position:** B2 **Answer:** $b^2 = a^2 + c^2$, obtained by matching the total arc length of the ellipse to the arc length of one period of the sine curve, then using that the ellipse's arc-length integral is strictly increasing in b^2.

Layer 1 THE WHAT — surface analysis

The natural attempt. "Rolling without slipping" is normally introduced as a pointwise, instantaneous condition: at every moment, the length of ellipse arc that has passed the contact

point equals the length of curve arc that has passed it. The honest first instinct is therefore to try to build that correspondence explicitly. Parametrise the ellipse by θ, $(a\cos\theta, b\sin\theta)$, so the ellipse arc-length element is $ds_{\text{ell}} = \sqrt{a^2\sin^2\theta + b^2\cos^2\theta}\,d\theta$; parametrise the curve by its own x-coordinate, with arc-length element $ds_{\text{curve}} = \sqrt{1 + (c/a)^2\cos^2(x/a)}\,dx$. Rolling without slipping says these must be equal at corresponding points, i.e. θ is a function of x satisfying the differential relation

$$\sqrt{a^2\sin^2\theta + b^2\cos^2\theta}\,\frac{d\theta}{dx} = \sqrt{1 + (c/a)^2\cos^2(x/a)}.$$

This looks like exactly the right thing to write down, and it is correct — but now try to solve it. The right-hand side, integrated, is an arc length of a sinusoid, which has no elementary antiderivative; the left-hand side, integrated in θ, is (up to scaling) the arc length of an ellipse, itself a non-elementary elliptic integral. There is no closed form for $\theta(x)$, and the problem gives no reason to expect one. A student who commits to solving for the pointwise correspondence has picked a genuinely unsolvable sub-problem, and can burn the rest of the exam discovering that fact.

The model chain. The problem does not ask for the contact point as a function of position; it gives exactly one piece of global data — one full revolution of the ellipse corresponds to exactly one period of the curve — and asks what that forces on a, b, c. So integrate the rolling condition over the whole range instead of trying to solve it pointwise.

The perimeter of the ellipse (integrating θ over a full turn) is

$$P_{\text{ell}} = \int_0^{2\pi} \sqrt{a^2\sin^2\theta + b^2\cos^2\theta}\,d\theta.$$

The arc length of the curve $y = c\sin(x/a)$ over one period, $x \in [0, 2\pi a]$, is

$$L_{\text{curve}} = \int_0^{2\pi a} \sqrt{1 + \left(\tfrac{c}{a}\cos\tfrac{x}{a}\right)^2}\,dx.$$

Substitute $\theta = x/a$, $dx = a\,d\theta$, θ ranging over $[0, 2\pi]$:

$$L_{\text{curve}} = \int_0^{2\pi} a\sqrt{1 + \left(\tfrac{c}{a}\cos\theta\right)^2}\,d\theta = \int_0^{2\pi} \sqrt{a^2 + c^2\cos^2\theta}\,d\theta.$$

Write $a^2 = a^2\sin^2\theta + a^2\cos^2\theta$ inside the square root:

$$L_{\text{curve}} = \int_0^{2\pi} \sqrt{a^2\sin^2\theta + (a^2 + c^2)\cos^2\theta}\,d\theta.$$

Rolling one full revolution per one period means $P_{\text{ell}} = L_{\text{curve}}$:

$$\int_0^{2\pi} \sqrt{a^2\sin^2\theta + b^2\cos^2\theta}\,d\theta = \int_0^{2\pi} \sqrt{a^2\sin^2\theta + (a^2 + c^2)\cos^2\theta}\,d\theta.$$

For fixed a and fixed θ, the integrand $\sqrt{a^2\sin^2\theta + B\cos^2\theta}$ is a strictly increasing function of $B \geq 0$ wherever $\cos^2\theta > 0$ (which holds for all but finitely many θ), so the definite integral, as a function of B, is strictly increasing on $B \geq 0$. Two values of B giving equal integrals must therefore be equal: $b^2 = a^2 + c^2$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same fact, which is not in dispute: rolling without slipping means the ellipse arc-length element and the curve arc-length element are equal at every corresponding instant, $ds_{\text{ell}} = ds_{\text{curve}}$. They split on what to do with that equation. The natural attempt treats it as a differential equation to be solved for the correspondence $\theta(x)$ itself. The model chain never solves for $\theta(x)$ at all — it integrates both sides over the one range the problem actually specifies (one full period / one full revolution) and compares only the two resulting numbers.

> *The problem tells me a global fact — one revolution happens over exactly one period — not a local one; it never asks where the contact point is at some particular x. So I don't need to know θ as a function of x anywhere in between the endpoints. I should integrate the rolling condition over the whole period first and see what single equation falls out, rather than trying to invert an equality of arc-length elements into an explicit correspondence that the problem never actually requires.*

Layer 3 THE WHY

Why is it legitimate to integrate first and skip the pointwise correspondence entirely, rather than this being a shortcut that happens to work but loses information? Because rolling without slipping, by its very definition, already guarantees that a well-defined, monotonic correspondence $\theta(x)$ exists: both $ds_{\text{ell}}/d\theta$ and ds_{curve}/dx are positive everywhere (they are arc-length elements of curves with nowhere-vanishing speed), so $d\theta/dx > 0$ throughout and θ increases continuously and monotonically as x does. The *only* extra content in "one revolution over one period" is a statement about the two endpoints of that already-guaranteed correspondence: θ goes from 0 to 2π exactly as x goes from 0 to $2\pi a$. Nothing about the interior of the motion is being specified beyond what the local rolling condition already forces automatically. So comparing the two total arc lengths captures every drop of information the hypothesis actually contains; trying to recover $\theta(x)$ in closed form would be solving for something the problem never made determinate in a usable way, and is not needed to answer the question asked.

The second load-bearing fact is why the integral equality can be legally converted into $b^2 = a^2 + c^2$ rather than left as an opaque relation between two integrals. This rests on $\sqrt{A + B\cos^2\theta}$ being strictly increasing in B for $B \geq 0$ (differentiate: $\partial_B \sqrt{A + B\cos^2\theta} = \cos^2\theta / (2\sqrt{A + B\cos^2\theta}) \geq 0$, with strict inequality off the measure-zero set $\cos\theta = 0$), so the definite integral over $\theta \in [0, 2\pi]$ inherits strict monotonicity in B (a sum, or limit of sums, of strictly-more-often-than-not strictly increasing terms is strictly increasing). This is what licenses solving an equation between two elliptic-type integrals *without ever evaluating either of them* — the comparison is done structurally, through monotonicity in the one free parameter, exactly the kind of move that becomes available only once the problem has been reduced to comparing totals rather than trajectories.

Caveat. The problem statement does not itself prove that "rolls without slipping, completing exactly one revolution per period" is realisable as an actual physical motion for every a, b, c satisfying $b^2 = a^2 + c^2$ (i.e. that a consistent, single-valued, monotonic $\theta(x)$ correspondence genuinely exists globally along the whole curve and not just matches at the two endpoints). The archive treats the total-arc-length argument as giving the necessary relation the problem asks for, which is what is required to answer "how are a, b, c related," but a full existence argument for the rolling motion itself is not constructed here.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "completes one revolution when it traverses one period" — a statement about totals at matching endpoints, not about the motion in between. On seeing it, write down immediately: this is an equation between two total arc lengths, not a request for the contact-point function.
2. That cue suggests setting up and equating the two definite integrals (ellipse perimeter, one period's arc length of the curve) rather than setting up a differential equation for the pointwise correspondence.
3. Hesitation occurs at the very first line, where "rolling without slipping" pulls toward the textbook instantaneous formulation $ds_{\text{ell}} = ds_{\text{curve}}$ as a relation to be solved pointwise. This is reinforced away by checking what data the problem actually supplies — a global count (one revolution per period), not a request to locate the contact point at a given x — before choosing local vs. global formulation.

4. The trap is chasing a closed form for $\theta(x)$, which mixes an elliptic integral (from the ellipse) against a sinusoidal arc-length integral (from the curve) and has no elementary solution; committing to it wastes the entire attempt on a sub-problem the question never required solving.

5. Recognise this family faster: whenever "rolling without slipping" is paired with a statement about a whole revolution, a whole period, or any other start-to-end count rather than a specific contact position, treat it as a total-arc-length equality problem immediately.

6. Reusable rule: before turning a stated instantaneous physical constraint into a differential equation, check whether the question only needs the constraint's integrated, boundary-to-boundary consequence — if so, integrate over the specified full range immediately and do not attempt the (possibly unsolvable) pointwise correspondence at all.

▷ **Cue:** two curves roll without slipping, and the given condition is a global count (one revolution per one period, or similar start-to-end matching) rather than a specific contact point $\longrightarrow$ **Move:** equate total arc lengths over the two matching global ranges instead of solving the local rolling condition for the pointwise correspondence

▷ **Cue:** an equation between two definite integrals of the form $\int \sqrt{A\sin^2\theta + B\cos^2\theta}\,d\theta$ over a common range, with one parameter B unknown $\longrightarrow$ **Move:** use that the integrand, hence the integral, is strictly increasing in $B \geq 0$, so equal integrals force equal B — no need to evaluate either integral in closed form

1995 B3

1995 B3. To each positive integer with n^2 decimal digits, we associate the determinant of the matrix obtained by writing the digits in order across the rows. For example, for $n = 2$, to the integer 8617 we associate $\det\begin{pmatrix} 8 & 6 \\ 1 & 7 \end{pmatrix} = 50$. Find, as a function of n, the sum of all the determinants associated with n^2-digit integers. (Leading digits are assumed to be nonzero; for example, for $n = 2$, there are 9000 determinants.)

Fields: linear algebra — multilinearity and alternation of the determinant **Position:** B3 **Answer:** The sum is 45 for $n = 1$, 20250 for $n = 2$, and 0 for every $n \geq 3$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The worked example already hands you the $n = 2$ method: the determinant $d_1d_4 - d_2d_3$ is linear in each row separately, so the sum over all 9000 matrices can be computed by multilinearity — sum over the first row, sum over the second row, and take the determinant of the two summed row-vectors. Carrying this out for $n = 2$ (only $d_1 \in \{1,\ldots,9\}$ is restricted, $d_2, d_3, d_4 \in \{0,\ldots,9\}$ freely) gives row-sums $R_1 = (450, 405)$, $R_2 = (450, 450)$, and $\det(R_1, R_2) = 450 \cdot 450 - 405 \cdot 450 = 20250$ — a genuine, checkable answer.

The natural next move is to try to push exactly this computation forward to general n: write the determinant via the Leibniz expansion, $\det M = \sum_{\sigma\in S_n} \operatorname{sgn}(\sigma) \prod_i d_{i,\sigma(i)}$, and sum this over all valid digit assignments term by term, tracking which digits are restricted (only $d_{1,1}$) and which range freely over 0–9. This is a real strategy — it is, after all, exactly what multilinearity is doing under the hood — but pursued directly through the $n!$-term expansion it turns into heavy bookkeeping for every n: for each permutation σ, the product $\prod_i d_{i,\sigma(i)}$ touches one digit from each row, and summing it over all $9 \cdot 10^{n^2-1}$ digit assignments requires carefully separating the one entry that's restricted to 1–9 (when $\sigma(1) = 1$) from all others. Nothing wrong happens, but nothing simplifies either — the calculation for $n = 3, 4, \ldots$ threatens to be exactly as long as the $n = 2$ case, times $n!$, with no sign that the answer collapses to 0 visible anywhere inside the arithmetic. A student who commits to grinding this out for $n = 3$ can spend the whole time budget re-deriving a fact (vanishing) that turns out to need no arithmetic at all.

The model chain. Generalise the $n = 2$ multilinearity trick itself, rather than the $n = 2$ Leibniz expansion. By iterated linearity of the determinant in each row (linear in row i for any fixed choice of the other $n - 1$ rows, for every i), summing over all valid digit matrices M gives

$$\sum_M \det M = \det(R_1, R_2, \ldots, R_n),$$

where R_i is simply the sum of row i's own vector over its own range of digit-tuples, independent of the other rows (this is exactly what produced $R_1 = (450, 405)$, $R_2 = (450, 450)$ for $n = 2$; it generalises verbatim to any number of rows, by pulling one row's sum out at a time).

Compute each R_i. Row $i \geq 2$ has all n entries ranging freely over $\{0, \ldots, 9\}$. Its j-th coordinate is

$$R_i[j] = \sum_{d_{i,j}=0}^{9} d_{i,j} \cdot 10^{n-1} = 45 \cdot 10^{n-1},$$

the same value for every coordinate j and — crucially — the same value for every row index $i \geq 2$, since every such row has exactly the same unrestricted range. So

$$R_2 = R_3 = \cdots = R_n = \left(45 \cdot 10^{n-1}, 45 \cdot 10^{n-1}, \ldots, 45 \cdot 10^{n-1}\right).$$

Row 1 differs only in its first coordinate, restricted to $\{1, \ldots, 9\}$: $R_1[1] = \sum_{d=1}^{9} d \cdot 10^{n-1} = 45 \cdot 10^{n-1}$ (same numeric value, coincidentally), and $R_1[j] = \sum_{d=0}^{9} d \cdot 9 \cdot 10^{n-2} = 405 \cdot 10^{n-2}$ for $j \geq 2$.

For $n \geq 3$ there are at least two rows among $R_2, \ldots, R_n$ (there are $n - 1 \geq 2$ of them), and they are literally identical vectors. A determinant with two equal rows is 0 (an immediate consequence of alternation: swapping those two identical rows must both negate the determinant and leave the matrix unchanged, forcing $\det = -\det$). Hence $\det(R_1, \ldots, R_n) = 0$ for all $n \geq 3$, with no further computation.

For $n = 2$ there is only one "other" row, so no repetition is available, and the answer is the direct computation already shown: $R_1 = (450, 405)$, $R_2 = (450, 450)$, $\det(R_1, R_2) = 450 \cdot (450 - 405) = 450 \cdot 45 = 20250$.

For $n = 1$ the "determinant" of a 1×1 matrix is the entry itself, and the sum over the nine possible leading digits $1, \ldots, 9$ is 45.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the identical tool, handed to the reader by the worked example: multilinearity converts a sum of determinants into the determinant of summed rows. They split on *what level to generalise at*. The natural attempt generalises the underlying Leibniz expansion — correct, but it re-derives the multilinearity trick from scratch inside a sum over permutations for every n, with no structural shortcut visible. The model chain generalises the trick itself, as a black box ('sum of dets = det of row-sums, for any number of rows'), and then asks a much smaller question: what do the row-sum vectors R_i look like, and are any two of them ever forced to coincide?

> *The $n = 2$ computation only needed one row-sum vector to be compared against another; but if I write down the row-sum trick for general n, most of the rows ($i \geq 2$) have exactly the same unrestricted range $\{0, \ldots, 9\}^n$ — so their row-sum vectors R_i don't just resemble each other, they are the* same *vector, for every $i \geq 2$. And a determinant with two identical rows is automatically zero, by the definition of what a determinant does under a row swap — I don't need to compute a single entry of R_1 to see this the moment $n \geq 3$ gives me two such repeated rows to work with.*

Layer 3 THE WHY

Why does noticing that $R_2 = R_3 = \cdots = R_n$ finish the problem outright, with no numerical work at all, once $n \geq 3$? Because the determinant is not merely linear in each row but *alternating*:

swapping two rows negates it. Apply that swap to two rows that happen to be identical, and the matrix itself is unchanged by the swap while the determinant is asserted to flip sign — so $\det = -\det$, forcing $\det = 0$. This is not a computational coincidence specific to this problem; it is a structural consequence of what a determinant *is* (an alternating multilinear form), and it is available the instant two of the summed row-vectors coincide, which happens here purely because rows 2 through n share the same digit range and thus the same row-sum — a fact about the problem's setup (only the very first digit of the whole number is restricted), not about any specific value of n.

This also explains, retroactively, why $n = 1$ and $n = 2$ are genuinely exceptional rather than small cases of a general pattern that happens to be nonzero: $n = 1$ has no "other" rows at all to repeat, and $n = 2$ has exactly one other row (R_2), so there is nothing for it to be identical *to* — the vanishing mechanism for $n \geq 3$ requires at least two unrestricted rows, which is precisely $n - 1 \geq 2$, i.e. $n \geq 3$.

Kedlaya's solution reaches the same conclusion for $n \geq 3$ by a different, also-structural route: it observes that permuting rows 2 through n of the underlying n^2-digit *integer* is a symmetry of the set of all n^2-digit integers with nonzero leading digit (since the restriction lives only in row 1), so summing over that set is unchanged by relabelling which row is which; but any single transposition of two rows negates the associated determinant. A sum that is invariant under an operation that negates every term must equal its own negative, hence 0. This is the alternation fact viewed from the outside, at the level of the summation index (permuting which matrices get summed) rather than from the inside, at the level of the row-sum vectors R_i — the two proofs are the same underlying fact (a determinant vanishes on repeated rows / is negated by a row swap) applied one derivative apart: to the individual terms' symmetry, or to the aggregated row-sums' equality. Neither is more correct than the other; the row-sum version is preferred here because it falls directly out of the multilinearity apparatus already built for $n = 2$, needing no second technique.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the worked $n = 2$ example itself using multilinearity (sum over rows, take determinant of the sums) to reach a checkable numeric answer. On seeing it, write down immediately that this technique is not specific to $n = 2$ — ask what it produces for general n before reaching for the Leibniz expansion.

2. That cue suggests generalising the row-sum trick as a black box: $\sum_M \det M = \det(R_1, \ldots, R_n)$ for any n, then computing only the row-sum vectors R_i rather than re-deriving the trick from the permutation expansion each time.

3. Hesitation occurs right after the $n = 2$ case succeeds: the instinct is to trust the method that just worked and push it forward exactly as executed (full Leibniz expansion, tracked digit by digit), rather than abstracting what made it work. This is reinforced away by asking "which part of this computation used $n = 2$ specifically, and which part is generic in n?" before starting the general computation.

4. The trap is spending effort re-verifying the multilinearity identity from scratch inside a permutation sum for every n, when the identity itself (sum of dets = det of row-sums) is already established once, generically, and the only remaining question — do any two row-sum vectors coincide? — is much smaller than the original problem.

5. Recognise this family faster: whenever a sum of determinants over independently-ranging rows appears, reach for the multilinearity identity $\sum \det M = \det(\text{row-sums})$ immediately, and specifically check whether any two rows share the exact same range of values — if they do, the sum vanishes with no arithmetic, by alternation.

6. Reusable rule: a determinant is an alternating multilinear form, so (a) sums over independently-ranging rows collapse to a single determinant of row-sums, and (b) that determinant vanishes for free the moment two of the summed rows are forced to be identical vectors — check for that symmetry before computing any entry.

- ▷ **Cue:** a quantity is defined as a sum of determinants over matrices whose rows range independently (some rows restricted, some not) ⟶ **Move:** use multilinearity to replace the sum by a single determinant of row-sum vectors, one vector per row, before attempting any term-by-term expansion
- ▷ **Cue:** two or more of the summed row-sum vectors (or, more generally, two rows/columns of a matrix) are forced to be identical by the problem's own symmetry ⟶ **Move:** conclude the determinant is 0 immediately, by alternation (a row swap both fixes an identical-row matrix and negates its determinant) — no further computation needed

1995 B4

1995 B4. Evaluate

$$\sqrt[8]{2207 - \frac{1}{2207 - \frac{1}{2207 - \cdots}}}.$$

Express your answer in the form $\frac{a+b\sqrt{c}}{d}$, where a, b, c, d are integers.

Fields: algebra; continued fractions and fixed points; the map $t \mapsto t + 1/t$ **Position:** B4 — mid-paper, but in practice one of the most accessible problems on the paper once the right invariant is seen; the cost of the wrong route is time, not failure **Answer:** $\frac{3+\sqrt{5}}{2}$, i.e. $a = 3, b = 1, c = 5, d = 2$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student writes L for the continued fraction, observes that the tail of the fraction is again L, and gets $L = 2207 - \frac{1}{L}$, that is

$$L^2 - 2207L + 1 = 0, \qquad L = \frac{2207 \pm \sqrt{2207^2 - 4}}{2}.$$

The value is large and positive, so take the $+$ sign. Now the arithmetic: $2207^2 - 4 = 4870845$. Factor it: $4870845 = (2207 - 2)(2207 + 2) = 2205 \cdot 2209 = (21^2 \cdot 5)(47^2)$, so $\sqrt{2207^2 - 4} = 987\sqrt{5}$ and

$$L = \frac{2207 + 987\sqrt{5}}{2}.$$

So far so good (and the factorisation $2207^2 - 4 = (2207 - 2)(2207 + 2)$ is itself a small win). Now the eighth root. The student guesses the answer has the form $\frac{p+q\sqrt{5}}{2}$ and wants $\left(\frac{p+q\sqrt{5}}{2}\right)^8 = \frac{2207+987\sqrt{5}}{2}$. Expanding an eighth power by the binomial theorem gives two polynomial equations of degree 8 in p, q — not solvable by inspection. The better version of this attempt denests three times: find $\frac{p+q\sqrt{5}}{2}$ whose square is L, i.e. $p^2 + 5q^2 = 2 \cdot 2207 = 4414$ and $pq = 987$. Now $987 = 3 \cdot 7 \cdot 47$; trying factor pairs, $p = 47, q = 21$ gives $47^2 + 5 \cdot 21^2 = 2209 + 2205 = 4414$. So $\sqrt{L} = \frac{47+21\sqrt{5}}{2}$. Repeat: $p^2 + 5q^2 = 94$, $pq = 21$ gives $p = 7, q = 3$. Repeat: $p^2 + 5q^2 = 14$, $pq = 3$ gives $p = 3, q = 1$. Answer $\frac{3+\sqrt{5}}{2}$. This works. It costs three rounds of Diophantine guessing, and it rests on an unproved assumption — that each square root has the form $\frac{p+q\sqrt{5}}{2}$ with integer p, q — which is justified only after the fact by finding the pair. It also silently assumes the continued fraction converges.

The model chain. **Step 1: the fraction converges, to the larger root.** The value is by definition the limit of $L_0 = 2207$, $L_{k+1} = f(L_k)$ with $f(t) = 2207 - \frac{1}{t}$. Let $r = \frac{2207+\sqrt{2207^2-4}}{2}$ be the larger root of $t^2 - 2207t + 1 = 0$, so $f(r) = r$ and $1 < r < 2207$. The map f is increasing on $t > 0$. Induction: $L_0 = 2207 > r$; if $L_k > r$ then $L_{k+1} = f(L_k) > f(r) = r$. And $L_1 = 2207 - \frac{1}{2207} < L_0$;

if $L_{k+1} < L_k$ then $L_{k+2} = f(L_{k+1}) < f(L_k) = L_{k+1}$. A decreasing sequence bounded below by r converges to some $L \geq r$ with $L = f(L)$, so L is a root of $t^2 - 2207t + 1$ that is $\geq r$: $L = r$. In particular

$$L + \frac{1}{L} = 2207, \qquad L > 1.$$

Step 2: halve the exponent three times. For any $u > 0$,

$$\left(\sqrt{u} + \frac{1}{\sqrt{u}}\right)^2 = u + \frac{1}{u} + 2.$$

So if $u + \frac{1}{u} = k$ then $\sqrt{u} + \frac{1}{\sqrt{u}} = \sqrt{k+2}$ (both sides positive). Apply this with $u = L, L^{1/2}, L^{1/4}$:

$$L^{1/2} + L^{-1/2} = \sqrt{2209} = 47, \quad L^{1/4} + L^{-1/4} = \sqrt{49} = 7, \quad L^{1/8} + L^{-1/8} = \sqrt{9} = 3.$$

Step 3: solve the last quadratic. $s = L^{1/8}$ satisfies $s^2 - 3s + 1 = 0$, so $s = \frac{3 \pm \sqrt{5}}{2}$. Since $L > 1$, $s > 1$, and $\frac{3-\sqrt{5}}{2} < 1$. Hence

$$\sqrt[8]{L} = \frac{3 + \sqrt{5}}{2}. \qquad \square$$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write $L = 2207 - \frac{1}{L}$. They split on what to do with that equation. (N) solves it — $L = \frac{2207 + 987\sqrt{5}}{2}$ — and then attacks the root of an explicit surd. (M) does not solve it; it rewrites it as $L + \frac{1}{L} = 2207$ and asks what the eighth root does to the quantity $t + \frac{1}{t}$. Once the invariant is $t + \frac{1}{t}$, taking a square root becomes "add 2, take a square root" on a single integer, and the three perfect squares $2209, 49, 9$ appear on their own.

> *I have $L^2 - 2207L + 1 = 0$. The constant term is 1, so the roots are L and $1/L$ and the equation says $L + \frac{1}{L} = 2207$. I don't want L; I want $L^{1/8}$. What does $t + \frac{1}{t}$ become when I replace t by $\sqrt{t}$? Its square is $t + \frac{1}{t} + 2$. So I add 2 and take a square root, three times.*

> *$2207 + 2 = 2209 = 47^2$. That is not an accident: the setter chose 2207 so that the chain would be perfect squares all the way down. If $k + 2$ is a square, halve again.*

- ▷ **Cue:** a quadratic $t^2 - kt + 1 = 0$, equivalently $t + \frac{1}{t} = k$, and a request for $t^{1/2^m}$ or t^{2^m} $\longrightarrow$ **Move:** track $t + \frac{1}{t}$: squaring sends $k \mapsto k^2 - 2$, square-rooting sends $k \mapsto \sqrt{k+2}$
- ▷ **Cue:** an infinite continued fraction or nested radical with a constant pattern $\longrightarrow$ **Move:** the value is a fixed point $L = f(L)$; prove convergence by monotonicity of f and pick the root the limit can reach
- ▷ **Cue:** a large, specific integer in the statement (2207) $\longrightarrow$ **Move:** test it against the operation the problem asks for: is $2207 + 2$ a square? Is it $k^2 - 2$ for small k? The number was chosen
- ▷ **Cue:** a request for $\sqrt[n]{\text{explicit surd}}$ $\longrightarrow$ **Move:** do not denest by guessing; find the polynomial the surd satisfies and transform the polynomial

Layer 3 THE WHY

Why is $t + \frac{1}{t}$ the natural thing to track, and not a trick? Go back to what the equation is. $L^2 - 2207L + 1 = 0$ has constant term 1, so by Vieta its two roots multiply to 1: they are L and $1/L$. Their sum is 2207. So the equation *is* the statement "$L + \frac{1}{L} = 2207$" — it carries no information about L beyond that symmetric function. Any question about L that is symmetric under $L \leftrightarrow 1/L$ should be answered in terms of $L + \frac{1}{L}$. And "the eighth root" respects that symmetry: $(1/L)^{1/8} = 1/L^{1/8}$. So the eighth root of the root-pair $\{L, 1/L\}$ is the root-pair

$\{s, 1/s\}$, and the whole problem lives on the one number $s + \frac{1}{s}$. That is why the model chain never needs to see $\sqrt{5}$ until the last line.

Why does square-rooting become "add 2 and square-root"? Because $(x + \frac{1}{x})^2 = x^2 + \frac{1}{x^2} + 2$ — the cross term is $2 \cdot x \cdot \frac{1}{x} = 2$, a constant, precisely because the product of the pair is 1. This is the same fact that gives $\cos 2\theta = 2\cos^2\theta - 1$: write $x = e^{i\theta}$ and $x + \frac{1}{x} = 2\cos\theta$. The map $k \mapsto k^2 - 2$ is the doubling map for anything of the form $x + x^{-1}$. It is a recursion worth owning: $k_0 = 3$, $k_{m+1} = k_m^2 - 2$ gives $3, 7, 47, 2207, \ldots$, and $k_m = \varphi^{2^{m+1}} + \varphi^{-2^{m+1}}$ with $\varphi = \frac{1+\sqrt{5}}{2}$ — these are the Lucas numbers $\ell_2, \ell_4, \ell_8, \ell_{16}$ (write $\ell_n = \varphi^n + (-1/\varphi)^n$, to keep them apart from the iterates L_k). That is also where the natural attempt's 987 came from: 987 is the Fibonacci number F_{16}, and $\varphi^{16} = \frac{\ell_{16} + F_{16}\sqrt{5}}{2}$. The student who recognised 2207 as a Lucas number would have been done in one line; the student who did not still gets there by the invariant, which is the point.

Why prove convergence? Because an infinite continued fraction is *defined* as a limit, and "the tail equals the whole" is only a statement about the limit if the limit exists. It also decides the sign: the equation has two positive roots, $r \approx 2207$ and $1/r \approx 0.00045$, and only monotonicity from $L_0 = 2207$ shows which one is reached. Kedlaya's solution says the sequence is "strictly decreasing (by induction)"; the induction is the one in Step 1, and it needs f increasing, which is where $f(t) = 2207 - \frac{1}{t}$ earns its shape.

Kedlaya then states: "if x satisfies $x^2 - ax + 1 = 0$ then $x^4 - (a^2 - 2)x^2 + 1 = 0$." Here is where that comes from: it is the doubling map $k \mapsto k^2 - 2$ read backwards, i.e. exactly Step 2. His product $(x^2 - ax + 1)(x^2 + ax + 1)$ is one way to derive it; the square $(x + \frac{1}{x})^2$ is shorter and shows why the constant is 2.

The natural attempt was not stupid — it got the right answer, and the factorisation $2207^2 - 4 = 2205 \cdot 2209$ is the same structure in disguise ($2209 = 47^2$ is the first rung of the ladder). It lost by solving the quadratic, which threw away the symmetric form and replaced one integer with a surd.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the equation $L^2 - 2207L + 1 = 0$, with constant term 1. On seeing it, write $L + \frac{1}{L} = 2207$.
2. The cue suggests working with the symmetric function $t + \frac{1}{t}$ under the operation requested (an eighth root, i.e. three square roots), using $(\sqrt{t} + 1/\sqrt{t})^2 = t + 1/t + 2$.
3. Hesitation occurs after obtaining $L = \frac{2207 + 987\sqrt{5}}{2}$ — the student sees the surd and feels committed to it. It is reinforced away by the habit: before solving a quadratic, ask what the problem wants from its roots, and whether that is symmetric.
4. Two traps. The first is denesting by guesswork, which works here but only because the setter made it work. The second is skipping convergence and the choice of root: $1/r$ is also a fixed point, and its eighth root is $\frac{3-\sqrt{5}}{2}$, also of the requested form.
5. Any time a problem involves x and $1/x$ together — a quadratic with constant term 1, a reciprocal equation, cos written as $\frac{z+z^{-1}}{2}$ — move to $x + \frac{1}{x}$ and ask how the requested operation acts on it.
6. Reusable decision rule: *when a problem asks for a function of a root of a polynomial, transform the polynomial, not the root. If the roots come in a pair $\{x, 1/x\}$, the invariant is $x + 1/x$, and powers act on it by $k \mapsto k^2 - 2$ (squaring) or $k \mapsto \sqrt{k+2}$ (square-rooting).*

1995 B5

1995 B5. A game starts with four heaps of beans, containing 3,4,5 and 6 beans. The two players move alternately. A move consists of taking **either**

a) one bean from a heap, provided at least two beans are left behind in that heap, **or**

b) a complete heap of two or three beans.

The player who takes the last heap wins. To win the game, do you want to move first or second? Give a winning strategy.

Fields: combinatorics; impartial games; invariants and parity **Position:** B5 — late in the paper by position, but it collapses once each heap is classified by what it can still do; Sprague–Grundy theory makes it routine, and a self-contained parity invariant makes it elementary **Answer:** Move first. Take one bean from the heap of 3, leaving heaps $2, 4, 5, 6$; thereafter always move so that the number of heaps of size 3 is even and the number of heaps whose size is 2 or an odd number ≥ 5 is even.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student first reads the rules carefully. From a heap of $h \geq 3$ you may go to $h-1$; a heap of 2 can only be removed; a heap of 3 can go to 2 or be removed. So a heap never reaches size 1, and the game ends exactly when every heap is gone: whoever makes the last move wins.

Then a good idea: count the moves. A heap of initial size $h \geq 3$ is whittled down one bean at a time until it is removed, either at size 3 (after $h-3$ single-bean moves, so $h-2$ moves in total) or at size 2 ($h-1$ moves in total). So

$$\#\text{moves in the whole game} = \sum_h (h-1) \;-\; k = 2+3+4+5-k = 14-k,$$

where k is the number of heaps that are removed at size 3. The first player wins iff $14-k$ is odd, i.e. iff k is odd. This is a real insight — the whole game is a fight over one parity. But now the student is stuck on the fight itself. Each heap, when it reaches 3, offers whoever is to move a choice: remove it (counts towards k) or reduce it to 2 (does not). Who gets that choice depends on the timing, and the timing depends on how many single-bean moves are available elsewhere, which both players can spend or save. The student tries "I'll make sure k is odd" and cannot see how to guarantee it against an opponent who stalls. The next resort is to search the game tree from $(3, 4, 5, 6)$: there are $4 \cdot 5 \cdot 6 \cdot 7$-ish positions, too many to do reliably by hand in the time.

The model chain. **Step 1: types.** Every heap has size 0 or ≥ 2. Classify a heap:

$$\text{type } 0: \text{ size 0 or even} \geq 4; \qquad \text{type } 1: \text{ size 2 or odd} \geq 5; \qquad \text{type } 2: \text{ size 3.}$$

List the moves on a single heap and the type change each causes:

even $h \geq 6 \to h-1$	$0 \to 1$
$4 \to 3$	$0 \to 2$
odd $h \geq 5 \to h-1$	$1 \to 0$
$2 \to 0$	$1 \to 0$
$3 \to 2$	$2 \to 1$
$3 \to 0$	$2 \to 0$

So every move changes the type of exactly one heap, to a different type.
Step 2: the invariant. Let n_1, n_2 be the numbers of heaps of type 1 and type 2. Call a position *balanced* if n_1 and n_2 are both even. Then:

(i) The final position (no heaps) is balanced.

(ii) Every move from a balanced position produces an unbalanced one. Indeed a single heap changes type $X \to Y$ with $X \neq Y$, so at least one of X, Y is 1 or 2. From the table: $0 \leftrightarrow 1$ flips the parity of n_1 only; $0 \leftrightarrow 2$ flips n_2 only; $2 \to 1$ flips both. In every case at least one of n_1, n_2 changes parity.

(iii) From every unbalanced position there is a move to a balanced one. If n_2 is odd and n_1 even, remove a heap of 3 ($2 \to 0$). If n_1 is odd and n_2 even, take a type-1 heap and move it to type 0 (remove it if it is 2, take one bean if it is odd ≥ 5). If both are odd, reduce a heap of 3 to 2 ($2 \to 1$). Each required heap exists because the corresponding count is odd, hence ≥ 1.

Step 3: the winner. The game is finite (every move removes at least one bean). If I hand my opponent a balanced position, then by (ii) every reply hands me an unbalanced one, and by (iii) I can hand back a balanced one. The only position with no moves is the empty one, which is balanced; so it is never my turn there, and it is my opponent who faces it — meaning I made the last move and took the last heap.
Step 4: the start. Heaps $3, 4, 5, 6$ have types $2, 0, 1, 0$: $n_1 = 1$, $n_2 = 1$, unbalanced. So the first player wins, with the move from case "both odd" of (iii): reduce the 3 to a 2. The position $2, 4, 5, 6$ has types $1, 0, 1, 0$: balanced. (It is also the *only* winning first move: $4 \to 3$, $5 \to 4$, $6 \to 5$ and $3 \to 0$ all leave one of n_1, n_2 odd.)
The strategy in words. After the first move, answer each opponent move as follows. If they take a bean from an even heap ≥ 6, take another bean from that same heap. If they reduce a 4 to a 3, remove that 3. If they remove a 2 or take a bean from an odd heap ≥ 5, do one of those two things to another heap of type 1. If they reduce a 3 to a 2, reduce another 3 to a 2. If they remove a 3, remove another 3. Each response restores the balance.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains read the rules the same way and both see that parity decides everything. They split at the unit of analysis. (N) counts moves over the *whole game* ($14 - k$) and then has to control a global quantity k against an opponent. (M) asks what each *single heap* is worth by itself — which sizes behave alike — and finds three classes. Once heaps are classified, a move is a change of one heap's class, and the invariant is a statement about two counts that can be checked move by move.

> *Heaps don't interact; a move touches one heap. So first I should understand one heap alone. From $h \geq 4$ there is only one move, $h \to h - 1$ — no choice until it hits 3. So only the parity of the distance to 3 matters, and a heap of 3 is the only place with a genuine choice.*

> *I want a set of positions that I can always return to and my opponent can never stay in, and that contains the empty position. Try: "evenly many heaps of each nonzero type." Check it against the move table.*

- ▷ **Cue:** a game made of independent heaps where each move touches one heap $\longrightarrow$ **Move:** classify single heaps first; find which sizes are interchangeable (Sprague–Grundy value), then combine
- ▷ **Cue:** a heap with only one legal move for a long stretch $\longrightarrow$ **Move:** only the parity of that stretch matters; collapse the stretch to its parity
- ▷ **Cue:** "who wins, and give a strategy" in a finite game $\longrightarrow$ **Move:** name the set P of positions you hand over; prove (i) the terminal position is in P, (ii) every move leaves P, (iii) every non-P position has a move into P
- ▷ **Cue:** counting total moves gives "constant $- k$" with k controlled by both players $\longrightarrow$ **Move:** the parity fight is local: find where the choices happen and treat each choice point as its own small game

Layer 3 THE WHY

Why classify single heaps? Because of the rules' own grammar: "a move consists of taking … from a heap." A move touches exactly one heap, and the legality of a move depends only

on that heap. So the game is literally a sum of four one-heap games, played by choosing a component each turn. When a game is a sum, the only honest first step is to understand the summands. The natural attempt's move count is correct, but it treats the game as one indivisible object and then has to reason about timing across heaps — which is exactly the interaction that the sum structure says does not exist.

Why these three classes? Look at one heap. From any $h \geq 4$ there is exactly one legal move, to $h-1$. A heap of 6 is therefore a heap of 5 with one extra forced move attached, and a heap of 5 is a heap of 4 with one extra move attached. Adding two forced moves to a component changes nothing in a sum game: whenever your opponent spends one, you spend the other. So $6 \sim 4$, $7 \sim 5$, and in general only the parity of $h-4$ matters for $h \geq 4$. A heap of 4 is a 3 with one forced move in front. A heap of 2 is exactly one move. A heap of 3 is the only choice point in the whole game: one move (remove) or two moves ($3 \to 2 \to 0$). In Sprague–Grundy language these are nim-values: $g(0) = 0$, $g(2) = \text{mex}\{0\} = 1$, $g(3) = \text{mex}\{1,0\} = 2$, $g(4) = \text{mex}\{2\} = 0$, $g(5) = \text{mex}\{0\} = 1$, $g(6) = \text{mex}\{1\} = 0$, alternating from there. My types are those values, and "balanced" is "the nim-sum $\bigoplus g = 0$": with values in $\{0,1,2\}$, the nim-sum's units bit is the parity of n_1 and its twos bit is the parity of n_2.

Kedlaya's solution says the problem "is dumb if you know the Sprague–Grundy theory". Here is where that comes from, and what to do if you do not know it: Sprague–Grundy is the statement that every one-heap game behaves in sums like a Nim heap of some size. You do not need the theorem to solve this problem; you need the three-part characterisation (i)–(iii) of the losing set, which is the definition of "the player to move loses" written as something checkable. Proving (i)–(iii) directly, as in Step 2, is a complete proof, and it takes a table and three cases.

Why does the move-count insight still matter? Because it explains the answer. The total number of moves is $14-k$; a 3-heap is the only place where k is decided, and the type-2 class is exactly the heaps that are, or will inevitably become, 3's at a moment when someone gets to choose. The first move $3 \to 2$ spends that choice immediately, in the first player's favour, and leaves a symmetric position ($2 \sim 5$, $4 \sim 6$ as pairs of equal type) that the first player can mirror. So the parity fight the natural attempt could not control is won by resolving the only live choice at once and then mirroring.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "a move consists of taking … from a heap": one heap per move. On seeing it, write the single-heap move graph: $h \to h-1$ for $h \geq 3$, $3 \to 0$, $2 \to 0$.

2. The cue suggests analysing one heap at a time (nim-values, or an ad hoc classification by parity) and combining heaps by counting parities of each class.

3. Hesitation occurs after computing "$14-k$ moves": the student knows what must be controlled but not how. It is reinforced away by switching from the global count to the local choice points — where, in one heap, does anyone ever have a choice?

4. The trap is the global move-count: correct, suggestive, and a dead end if pursued as a timing argument. A second trap is misreading the rules — a heap can never be reduced to 1, which is what makes the empty position the only terminal one and "last heap" the same as "last move".

5. Whenever a game is played on several independent components, one component per move, ask first what one component is worth; long forced stretches collapse to their parity, and choice points are the only real content.

6. Reusable decision rule: *to solve a finite two-player game, guess a set P of positions (from small cases or by classifying components) and verify three things: the terminal position is in P, no move stays in P, and every position outside P has a move into P. The winning strategy is "always move into P".*

1995 B6

1995 B6. For a positive real number α, define

$$S(\alpha) = \{\lfloor n\alpha \rfloor : n = 1, 2, 3, \dots\}.$$

Prove that $\{1, 2, 3, \dots\}$ cannot be expressed as the disjoint union of three sets $S(\alpha), S(\beta)$ and $S(\gamma)$. [As usual, $\lfloor x \rfloor$ is the greatest integer $\leq x$.]

Fields: number theory; Beatty sequences; fractional parts and Diophantine approximation **Position:** B6 — the last problem, and genuinely hard: the official solution uses two-dimensional Weyl equidistribution. It yields to an elementary argument (a counting identity plus Dirichlet's pigeonhole approximation), but finding that argument is the whole difficulty **Answer:** No such partition exists. If it did, $F(M) = \{-M/\alpha\} + \{-M/\beta\} + \{-M/\gamma\}$ would equal 2 for every $M \geq 1$; choosing N with $N/\alpha, N/\beta, N/\gamma$ all very close to integers forces $F(N-1) = 1$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student knows Beatty's theorem: for irrational $\alpha, \beta > 1$ with $\frac{1}{\alpha} + \frac{1}{\beta} = 1$, the sets $S(\alpha), S(\beta)$ partition the positive integers. So two Beatty sets can do it; the problem says three cannot. The student starts with what is forced.

First, $\alpha, \beta, \gamma > 1$: if $\alpha < 1$ then $\lfloor \alpha \rfloor = 0 \in S(\alpha)$, which is not a positive integer; if $\alpha = 1$ then $S(\alpha)$ is everything and leaves no room. Second, density. Since $\alpha > 1$, $n \mapsto \lfloor n\alpha \rfloor$ is injective, and $S(\alpha)$ has about K/α elements up to K. The three sets partition, so $\frac{1}{\alpha} + \frac{1}{\beta} + \frac{1}{\gamma} = 1$. Third, rationality: if $\alpha = p/q$ then $S(\alpha) \ni \lfloor kq \cdot p/q \rfloor = kp$ for every k, so $S(\alpha)$ contains all multiples of p. Two rational parameters p/q, r/s would both contain pr. So at most one is rational.

Now the student is at the edge of what is standard. One idea: the complement of $S(\alpha)$ (for irrational α) is $S(\alpha')$ with $\frac{1}{\alpha'} = 1 - \frac{1}{\alpha}$, so $S(\beta) \cup S(\gamma) = S(\alpha')$, and one must show a Beatty set cannot be split into two Beatty sets. That is a restatement, not progress. Another idea: characterise membership. $m \in S(\alpha)$ iff some n has $m \leq n\alpha < m + 1$, iff the interval $[\frac{m}{\alpha}, \frac{m+1}{\alpha})$ contains an integer; for irrational α this is iff $\{\frac{m+1}{\alpha}\} < \frac{1}{\alpha}$. So the partition says: for every m, exactly one of three fractional-part conditions holds. To get a contradiction one needs an m where zero or two hold, i.e. control of $\{\frac{m}{\alpha}\}$ and $\{\frac{m}{\beta}\}$ *simultaneously*. The student may know that $\{n\theta\}$ is dense for irrational θ, but density of each coordinate separately says nothing about the pair, and the student stalls. (Kedlaya's solution gets past this wall with the two-dimensional Weyl equidistribution theorem — and then needs a separate argument for the case where $1, \frac{1}{\alpha}, \frac{1}{\beta}$ are rationally dependent.)

The model chain. Write $x = \frac{1}{\alpha}, y = \frac{1}{\beta}, z = \frac{1}{\gamma}$. As above, $\alpha, \beta, \gamma > 1$, so $x, y, z \in (0, 1)$.

Step 1: exact count. For $t \in (0, 1)$ let $c_t(K) = \#\big(S(1/t) \cap \{1, \dots, K\}\big)$. Every element $\lfloor n/t \rfloor$ is $\geq \lfloor 1/t \rfloor \geq 1$, and $n \mapsto \lfloor n/t \rfloor$ is injective because consecutive values of n/t differ by $1/t > 1$. And $\lfloor n/t \rfloor \leq K \iff n/t < K + 1 \iff n < (K+1)t$. So

$$c_t(K) = \#\{n \geq 1 : \ n < (K+1)t\} = \lceil (K+1)t \rceil - 1.$$

The partition gives $c_x(K) + c_y(K) + c_z(K) = K$ for every $K \geq 0$. Put $M = K + 1 \geq 1$:

$$\lceil Mx \rceil + \lceil My \rceil + \lceil Mz \rceil = M + 2 \qquad (M \geq 1). \tag{1}$$

Dividing by M and letting $M \to \infty$ gives $x + y + z = 1$.

Step 2: the fractional-part identity. For real u, $\lceil u \rceil = u + \{-u\}$. Substituting in (1) and using $Mx + My + Mz = M$:

$$F(M) := \{-Mx\} + \{-My\} + \{-Mz\} = 2 \qquad \text{for every integer } M \geq 1. \tag{2}$$

(Check at $M = 1$: $F(1) = (1 - x) + (1 - y) + (1 - z) = 2$. Consistent, as it must be.)

Step 3: an N that is nearly a common period. Let $\mu = \min\{x, 1-x, y, 1-y, z, 1-z\} > 0$ and choose an integer Q with $2/Q < \mu$. Consider the Q^2+1 points $(\{nx\}, \{ny\})$, $n = 0, 1, \ldots, Q^2$, in $[0,1)^2$, cut into Q^2 half-open squares of side $1/Q$. Two of them, $n_1 < n_2$, share a square. Put $N = n_2 - n_1 \geq 1$. Then

$$Nx = p_x + e_x, \qquad Ny = p_y + e_y, \qquad p_x, p_y \in \mathbb{Z}, \quad |e_x|, |e_y| < \tfrac{1}{Q},$$

with $e_x = \{n_2 x\} - \{n_1 x\}$ and $p_x = \lfloor n_2 x \rfloor - \lfloor n_1 x \rfloor$ (same for y). Set $p_z = N - p_x - p_y$ and $e_z = Nz - p_z$. Since $x + y + z = 1$, $e_z = -(e_x + e_y)$, so $|e_z| < 2/Q$. Thus

$$e_x + e_y + e_z = 0, \qquad |e_t| < \mu \leq \min(t, 1-t) \quad (t = x, y, z).$$

Also $N \geq 2$: if $N = 1$ then $|x - p_x| < \min(x, 1-x)$, but the distance from $x \in (0,1)$ to the nearest integer is exactly $\min(x, 1-x)$.

Step 4: evaluate $F(N-1)$. For $t \in \{x, y, z\}$,

$$-(N-1)t = -p_t - e_t + t, \qquad \text{so} \qquad \{-(N-1)t\} = \{t - e_t\} = t - e_t,$$

because $|e_t| < \min(t, 1-t)$ gives $t - e_t > t - t = 0$ and $t - e_t < t + (1-t) = 1$. Summing,

$$F(N-1) = (x + y + z) - (e_x + e_y + e_z) = 1 - 0 = 1.$$

But $N - 1 \geq 1$, so (2) says $F(N-1) = 2$. Contradiction. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the density relation $x + y + z = 1$, and both translate membership into fractional parts. They split at *which* fractional-part statement to aim at. (N) works set by set — "$m \in S(\alpha)$ iff $\{\frac{m+1}{\alpha}\} < \frac{1}{\alpha}$" — and then needs to control two or three fractional parts at once in a two-dimensional region, which is where Weyl's theorem is invoked. (M) sums the three counts first, so the partition becomes a single identity $F(M) = 2$ for all $M \geq 1$, and then asks what that identity would say at a *negative* argument. Since $F(-1) = \{x\} + \{y\} + \{z\} = 1$, not 2, all that is needed is a positive M that behaves like -1, and Dirichlet's pigeonhole principle supplies one: $M = N - 1$ with N an approximate common period.

> *The partition is really a statement about counts: the three counting functions add up to K for every K. Write that exactly with ceilings, subtract off the linear part, and it says $\{-Mx\} + \{-My\} + \{-Mz\} = 2$ for every positive M. That's one clean identity; attack it, not the sets.*

> *What would the identity say for $M = -1$? $\{x\} + \{y\} + \{z\} = x + y + z = 1$. Not 2. The fractional parts flip under $M \mapsto -M$: $\{u\} + \{-u\} = 1$. So I need a positive integer that looks like -1 to all three numbers at once — $N - 1$, where Nx, Ny, Nz are all almost integers. Pigeonhole gives such an N.*

- ▷ **Cue:** a partition of $\mathbb{N}$ into sets whose counting functions are known ⟶ **Move:** add the counting functions: the sum equals K exactly, for every K; this is an identity, not an estimate
- ▷ **Cue:** $\lceil u \rceil$ or $\lfloor u \rfloor$ summed and compared with a linear quantity ⟶ **Move:** write $\lceil u \rceil = u + \{-u\}$; the linear parts cancel and a fractional-part identity remains
- ▷ **Cue:** an identity $F(M) = c$ for all $M \geq 1$ with F built from $\{Mt\}$ ⟶ **Move:** test it at $M = -1$ (using $\{u\} + \{-u\} = 1$), then transport that value to a positive $M = N - 1$ with N an approximate common period
- ▷ **Cue:** need $Nt_1, \ldots, Nt_k$ all within ε of integers ⟶ **Move:** Dirichlet's pigeonhole: $Q^k + 1$ points $(\{nt_i\})$ in Q^k boxes; the difference of two indices in the same box works

Layer 3 THE WHY

Why sum the counts? Because "$\{1,2,3,\dots\}$ is the disjoint union" is, restricted to $\{1,\dots,K\}$, the statement that the three counts add to K — and for $\lfloor n\alpha \rfloor$ the count is an explicit ceiling. Each set separately carries messy information (which m are in it); the sum carries exact information (the counts add to K, with no error term). An exact identity for all K is far stronger than the asymptotic density statement that most solvers extract from it and then discard. The natural attempt took the limit, got $x+y+z=1$, and threw away the identity it came from. The model keeps the identity.

Why look at $M=-1$? Because F is built from fractional parts, and the one structural fact about fractional parts under negation is $\{u\}+\{-u\}=1$ for $u \notin \mathbb{Z}$. So for any M with all Mt non-integral, $F(-M)=3-F(M)$. If F were given by a nice formula on all of $\mathbb{Z}$, it could only be constant if that constant were $\frac{3}{2}$ — impossible for an integer. That is the heart of the problem: an identity with value 2 on positive integers "wants" value 1 on negative ones. For *two* Beatty sets the same computation gives $F \equiv 1$ on positives and $2-1=1$ on negatives: no conflict, which is exactly why Beatty's theorem is true. For k sets, $F \equiv k-1$ on positives and 1 on negatives: a conflict for every $k \geq 3$. The problem is really about this asymmetry.

Why is $N-1$ allowed to stand in for -1? Because F depends on M only through the fractional parts of Mx, My, Mz, and those are unchanged to within ε when M is shifted by a common near-period N. We only need approximate periodicity, not exact, because $F(N-1)$ is computed exactly once the errors are smaller than the distances of x,y,z from the integers — that is what the condition $|e_t| < \min(t, 1-t)$ buys. And we need the three errors to sum to exactly 0, which is free: $x+y+z=1$ and N is an integer.

Why pigeonhole and not Weyl? Weyl's theorem says the points $(\{nx\},\{ny\})$ fill the square evenly; that is far more than needed and it fails when $1, x, y$ are rationally dependent, which forces Kedlaya into a second case. Dirichlet's argument only asks for *one* n with both coordinates near 0, and pigeonhole gives that for any reals, rational or not. Note that the model never needed "at most one of α,β,γ is rational" — the rational cases are covered automatically. Kedlaya's solution uses the equidistribution to find n with $\{-(n+1)/\alpha\}+\{-(n+1)/\beta\}<1$. Here is where that comes from: it is precisely $F(n+1)\leq 1$, the same contradiction, and the model finds that $n+1$ as $N-1$ instead of searching the whole square for it.

A note on provenance: this elementary proof is not the official route. It was constructed for the archive and checked step by step (and numerically: for random $x+y+z=1$, F takes both values 1 and 2 on positive integers, as the argument predicts it must). The same argument with k sets and Dirichlet approximation of $k-1$ reciprocals gives Uspensky's theorem: no partition of the positive integers into $k \geq 3$ Beatty sets.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "disjoint union" of sets with explicit counting functions. On seeing it, write $\sum_t \left(\lceil (K+1)t \rceil - 1\right) = K$ for all K.

2. The cue suggests converting the counting identity into a fractional-part identity $F(M)=2$ and seeking an M that violates it — via Diophantine approximation of the three reciprocals.

3. Hesitation occurs after $x+y+z=1$: the density fact is true but useless, and the solver does not know what to do next. It is reinforced away by refusing to take the limit: keep the exact identity and ask what values F can take.

4. The trap is to reach for equidistribution in two dimensions and then drown in the rational-dependence cases. A second trap is to treat the problem as three separate membership conditions rather than one summed identity.

5. When an identity involving $\{Mt\}$ holds for all positive M, evaluate it at $M=-1$ or $M=-m$; if the value disagrees, pull that negative argument into the positives with an approximate common period from pigeonhole.

6. Reusable decision rule: *an exact identity for all $M \geq 1$ is worth more than its limit — keep it. For fractional-part identities, use $\{u\}+\{-u\}=1$ to find the "mirror" value at negative M, and*

Dirichlet's simultaneous approximation to realise a negative M as $N - M$ with N positive.